Monitoring-based Diagnosis of Discrete-Event Systems with Uncertain Observations

Gianfranco Lamperti Marina Zanella

Dipartimento di Elettronica per l'Automazione
Università di Brescia, Italy
lamperti@ing.unibs.it zanella@ing.unibs.it

Abstract

A technique for diagnosis of a class of asynchronous discrete-event systems is presented. It copes with uncertain observations while monitoring the system, by generating diagnostic information at the occurrence of each new fragment of observation. Uncertainty may stem from noise affecting the communication channels, and from the multiplicity of such channels, which is bound to relax the absolute temporal ordering of the observable events. The challenge in large-scale application domains is twofold: reasoning without any explicit global model of the system, and incrementally generating the knowledge structures, taking into account that estimates of the system state and the relevant candidate diagnoses may not survive the occurrence of a new piece of observation.

1 Introduction

Diagnosis is the task of finding out the faults affecting a system based on observed symptoms. Although a central topic in Artificial Intelligence (AI) and historically one of the first to be tackled, automated diagnosis is still a research subject. Up to the middle of the 1990s, diagnostic reasoning mainly focused on static and quasi-static systems, while in the last decade the applicability of model-based diagnosis to the larger class of *dynamic systems* has been investigated [Struss, 1997; Brusoni *et al.*, 1998; Console *et al.*, 2002]. This led to the awareness that such a task depends on the ability to estimate the state based on observations.

Discrete-event systems (DESs) are a qualitative abstraction of continuous dynamic systems that has been receiving increasing attention from the AI community [Rozé and Cordier, 2002; Lamperti and Zanella, 2003; Grastien *et al.*, 2004; Schumann *et al.*, 2004]. Each state variable of a DES can only range over a finite number of symbolic values and the behavior of the DES can be described by means of state changes driven by a finite set of events. In the literature (as well as in this paper) the topology of each DES is usually distributed [Pencolé, 2004], specified as a network of components. The behavioral model of each component is a communicating automaton representing a nondeterministic and complete behavior. The challenge with large-scale DESs is reasoning without any explicit behavioral model of the whole system.

This paper deals with diagnosis while monitoring a DES. A diagnostic result is produced by processing each newly occurred observation fragment. At each step, the reasoning engine starts from the current system state and finds out all possible evolutions (sequences of transitions) that comply with the given fragment, thus reaching the next system state. However, since there may be several evolutions complying with the same observable event (especially when such an event is uncertain), each estimated system state is a hyperstate (including several candidate system states). Moreover there may be system states in the current hyperstate that will not survive the next processing step. At each step, the incremental problem-solving method can either throw away all the hyperstates apart from the current one (thus reducing space) or keep them in memory so as to reuse them (thus increasing efficiency). The approach proposed in this paper is inspired by the *bridged diagnostic method* [Lamperti and Zanella, 2004a] and by the notion of an uncertain observation [Lamperti and Zanella, 2002]. The latter allows the modeling of real world observations that are uncertain as to the identity of the observed labels (logical uncertainty) and/or their reciprocal emission order (temporal uncertainty) and/or the identity of the component emitting the label (source uncertainty). The main difference between an uncertain and a fragmented observation is that the former cumulatively represents all the observable events received over a time interval (and therefore it is an input for *a posteriori diagnosis*, which is the only task faced in [Lamperti and Zanella, 2002]) while the latter represents a single uncertain observable event, namely a *message*. No previous contribution in the literature to monitoring-based diagnosis of DESs considers uncertain observable events.

2 Modeling

A *system* is a network of *components* connected to one another through *links*. Each component is modeled by a communicating automaton that reacts to events either from the external world or from neighboring components through links. Formally, the automaton is a 6-tuple $(\mathbf{S}, \mathbf{E}_{\text{in}}, \mathbf{I}, \mathbf{E}_{\text{out}}, \mathbf{O}, \mathbf{T})$, where $\mathbf{S}$ is the set of states, $\mathbf{E}_{\text{in}}$ the set of input events, $\mathbf{I}$ the set of input terminals, $\mathbf{E}_{\text{out}}$ the set of output events, $\mathbf{O}$ the set of output terminals, and $\mathbf{T}$ the nondeterministic transition function $\mathbf{T} : \mathbf{S} \times \mathbf{E}_{\text{in}} \times \mathbf{I} \times 2^{\mathbf{E}_{\text{out}} \times \mathbf{O}} \mapsto 2^{\mathbf{S}}$. A transition $T \in \mathbf{T}$, from S to S', that is triggered by event e at input terminal I, and generates events $e_1, \ldots, e_k$ at output terminals $O_1, \ldots, O_k$, respectively, is denoted by

$$T = S \xrightarrow[(e_1,O_1),\ldots,(e_k,O_k)]{(e,I)} S'.$$

An In terminal is implicitly assumed, that is meant to receive events from the external environment. Links store the events

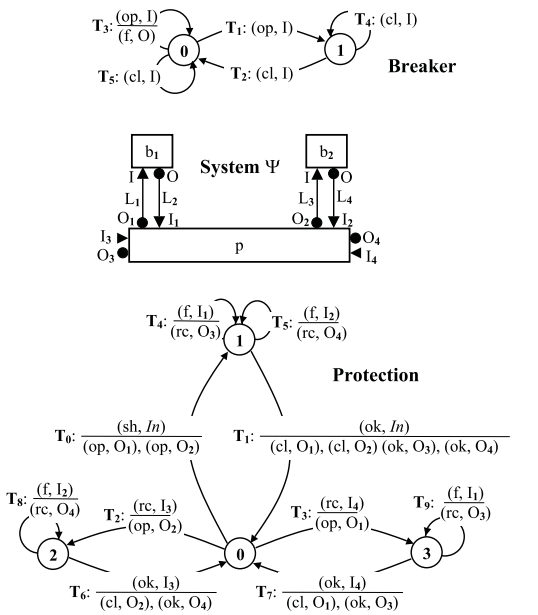

Figure 1: System Ψ and relevant component models.

exchanged between components. A link is characterized by a *management policy*, which establishes the effect of each event insertion or consumption. In this paper we confine the management policy to a FIFO queue. This way, the only significant parameter of the link is its *capacity*, the maximum number of events that can be buffered. The insertion of an event into a full link results in the loss of the event.

Example 1. Centered in Fig. 1 is a system Ψ made of protection p and breakers b_1 and b_2. Triangles and bullets denote input and output terminals, respectively. Breakers are connected with p through links $L_1 \cdots L_4$, while p is assumed to receive/send events from/to adjacent systems via *dangling* terminals I_3, I_4 / O_3, O_4. Ψ is an abstraction of the protection apparatus of a power transmission line, where breakers are tripped by p when a short circuit occurs on the line. The breaker model (top of Fig. 1) involves states 0 (closed) and 1 (open). Transitions are triggered by events *op* and *cl* received at terminal I. In T_3, the breaker cannot open: this is signaled to p by output event f. The protection model (bottom of Fig. 1) embodies states 0 (line is normal), 1 (line is shorted), 2 (recovery action requested by left-hand side line), and 3 (recovery action requested by right-hand side line). A recovery action is needed when a breaker fails to open (this causes the enlargement of the isolation). Event *ok* (the short has extinguished) moves p to normal state anew.

3 Behavior

Given an initial state Σ_0, system Σ evolves in a way that is both consistent with its topology and component models. The graph representing the whole set of possible evolutions is the *domain* of Σ *rooted* in Σ_0, $Dom(\Sigma, \Sigma_0)$. Such a graph is connected since an assumption of the approach is that the initial state of the system to be monitored is univocally known. The *behavior* $Bhv(\Sigma, \Sigma_0)$ is a connected subgraph of $Dom(\Sigma, \Sigma_0)$, which is restricted by *domain-dependent constraints*. A (possibly empty) path between two nodes of the behavior is a *history segment* of Σ. In particular, if the starting node is the initial state of the behavior, such a path is a *history* of Σ.

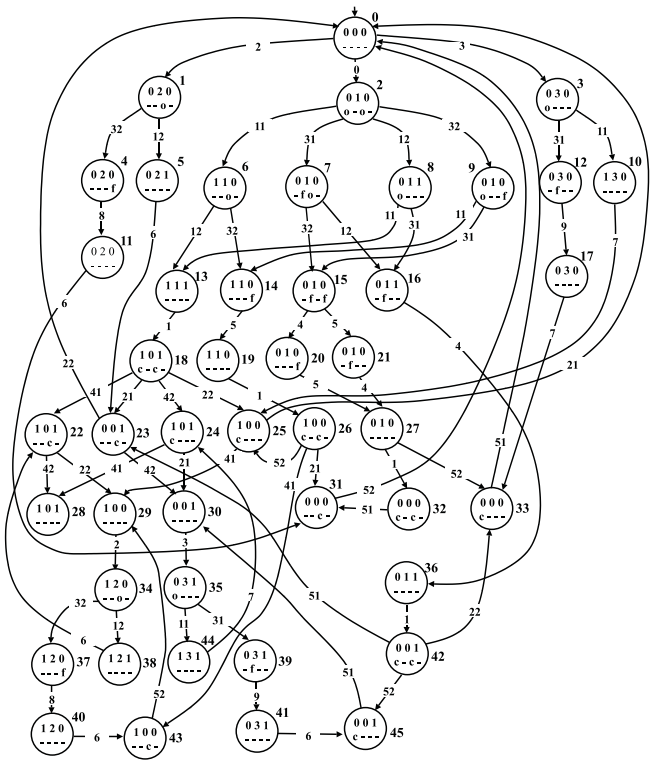

Figure 2: Behavior $Bhv(\Psi, \Psi_0)$.

Example 2. Shown in Fig. 2 is a $Bhv(\Psi, \Psi_0)$, where $\Psi_0 = (\mathbb{S}_0, \mathbb{L}_0)$, $\mathbb{S}_0 = (0, 0, 0)$, and $\mathbb{L}_0$ involves empty links. In each node, the record $\mathbb{S}$ of the component states for b_1, p, and b_2 is on the top, while the record $\mathbb{L}$ of queues of events within links $L_1 \cdots L_4$ is on the bottom. Incidentally, at most one event is stored in each link, so that the queue in the link can be expressed by either the event or a dash (empty link). Labels o and c are a shorthand for *op* and *cl*, respectively. Each edge is marked by a label identifying a component transition: single digits refer to transitions of the protection, while two-digit strings correspond to breaker transitions. For example, 2, 9, 31, and 42 stand for $T_2(p)$, $T_9(p)$, $T_3(b_1)$, and $T_4(b_2)$, respectively. A history $h(\Psi)$ is identified by the sequence of labels (component transitions) marking its edges, as for instance, $h(\Psi) = \langle 0, 31, 12, 4 \rangle$, corresponding to the following scenario: (i) a short circuit occurs on the line protected by Ψ and p commands both b_1 and b_2 to open, (ii) b_1 fails to open, while (iii) b_2 opens correctly, and (iv) p asks the left-hand side protection a recovery action. □

The concept of behavior has been introduced in this section for presentation purposes only. A behavior is a piece of knowledge that is implicit in the models inherent to the system at hand, that is, the topology of the system, the models of its components, and the domain-dependent constraints. The problem solving method for monitoring-based diagnosis proposed in this paper (Section 5), however, does not need any explicit behavior to be drawn from such models, the same as all the methods in previous contributions by the authors.

4 Diagnostic Problem

A diagnostic problem concerns the system model, the observer of the system evolution, the observation, and the characterization of faulty behavior. Let $\mathbf{T}$ be the set of transitions of components in Σ, and $\mathbf{V}$ a domain of labels including the null label ε. A *viewer* $\mathcal{V}$ is a mapping from $\mathbf{T}$ to $\mathbf{V}$. If $(T, \varepsilon) \in \mathcal{V}$ then T is *silent* else T is *visible*. The *product* of a history h and $\mathcal{V}$ is the sequence of labels $h \boxtimes \mathcal{V} = \langle \ell \mid T \in h, (T, \ell) \in \mathcal{V}, \ell \neq \varepsilon \rangle$.

Example 3. A viewer $\mathcal{V}_\psi$ for Ψ can be defined by the set of visible transitions, $\mathcal{V}_\psi = \{(T_0(p), sh), (T_2(p), l), (T_3(p), r), (T_1(b_1), o_1), (T_2(b_1), c_1), (T_1(b_2), o_2), (T_2(b_2), c_2)\}$. □

When the system is operating, each visible transition is perceived by $\mathcal{V}$ as a *message*. Each message μ is a pair (λ, τ), where $\lambda = \{\ell_1, \ldots, \ell_k\}$ is a subset of $\mathbf{V}$, namely the *logical content*, while $\tau = \{\mu'_1, \ldots, \mu'_h\}$ is a set of messages, namely the *temporal content*, identifying all the messages temporally preceding the current one. A *fragmented observation* is a list of messages, $\mathcal{O} = \langle \mu_1, \ldots, \mu_n \rangle$, where the following is assumed: $\forall i \in [1\,..\,n], \mu_i = (\lambda_i, \tau_i)\,(\tau_i \subseteq \{\mu_1, \ldots, \mu_{i-1}\})$. A message is uncertain, both logically and temporally. *Logical uncertainty* means that λ includes the actual label associated with the transition that generated it, but further *spurious* labels may be involved too. *Temporal uncertainty* means that only partial ordering is known among messages. The assumption making an integral part within the definition of a fragmented observation does not prevent time switching between emission and reception of any two messages. What it states is that, if a message a, emitted before message b, is received after b, it is impossible to the viewer to know that a precedes b, that is, the relative emission order of the two messages is unknown to the viewer. A fragmented observation may be mapped to an *observation graph* $\gamma(\mathcal{O}) = (\Omega, \Upsilon)$, a DAG where Ω is the set of nodes isomorphic to the messages in $\mathcal{O}$ and Υ the set of edges isomorphic to the temporal contents of messages. A *sub-observation* $\mathcal{O}_{[i]}$ of $\mathcal{O}$, $i \in [0\,..\,n]$, is the (possibly empty) prefix of $\mathcal{O}$ up to the i-th message, $\mathcal{O}_{[i]} = \langle \mu_1, \ldots, \mu_i \rangle$. When $\mu_1 = (\{\ell_1\}, \emptyset)$ and $\forall i \in [2\,..\,n]\,(\mu_i = (\{\ell_i\}, \{\mu_{i-1}\}))$, $\mathcal{O}$ is a *plain observation*, and is denoted by the list $\langle \ell_1, \ldots, \ell_n \rangle$ of *plain messages*.

Example 4. A fragmented observation relevant to viewer $\mathcal{V}_\psi$ is $\mathcal{O}_\psi = \langle \mu_1, \ldots, \mu_6 \rangle$, where $\mu_1 = (\{sh\}, \emptyset)$, $\mu_2 = (\{o_1\}, \{\mu_1\})$, $\mu_3 = (\{l, \varepsilon\}, \{\mu_1\})$, $\mu_4 = (\{o_2\}, \{\mu_2, \mu_3\})$, $\mu_5 = (\{c_2\}, \{\mu_4\})$, and $\mu_6 = (\{c_1, r\}, \{\mu_5\})$. The relevant observation graph $\gamma(\mathcal{O}_\psi)$ is depicted on the left of Fig. 3. □

Since it is neither trivial nor efficient to reason about the observation graph as is, an additional (acyclic) automaton is considered, called the *index space* of the observation, $\mathfrak{I}(\mathcal{O}) = (\mathbb{S}, \mathbb{E}, \mathbb{T}, S_0, \mathbb{S}_\mathrm{f})$, where $\mathbb{S}$ is the set of states, $\mathbb{E} = \mathbf{V} - \{\varepsilon\}$ the set of events, $\mathbb{T}$ the transition function, S_0 the initial state, and $\mathbb{S}_\mathrm{f}$ the set of final states. In $\mathfrak{I}(\mathcal{O})$ each path from the root to a final node, called a *temporal sequence*, represents a mode in which labels may be chosen in the observation graph without violating the constraints imposed by temporal and logical uncertainty. The whole set of such paths is the *extension* of $\mathfrak{I}(\mathcal{O})$, denoted $\|\mathfrak{I}(\mathcal{O})\|$.

Example 5. Shown on the right of Fig. 3 is the index space $\mathfrak{I}(\mathcal{O}_\psi)$, whose extension $\|\mathfrak{I}(\mathcal{O}_\psi)\|$ includes six temporal sequences. $\Im_7$ is the only final state. □

Let $\mathbf{T}$ be the set of transitions relevant to components in Σ, and $\mathbf{R}$ a set of labels including the *null* label ε. A *ruler* $\mathcal{R}$ is a mapping from $\mathbf{T}$ to $\mathbf{R}$. If $(T, \varepsilon) \in \mathcal{R}$ then T is *normal* else T is *faulty*. A subset $\delta \subseteq (\mathbf{R} - \{\varepsilon\})$ is a *generic diagnosis*. The *product* of a history segment h and $\mathcal{R}$ is a generic diagnosis $h \otimes \mathcal{R} = \{\varphi \mid T \in h, (T, \varphi) \in \mathcal{R}, \varphi \neq \varepsilon\}$. A set of generic diagnoses is a *diagnostic set*.

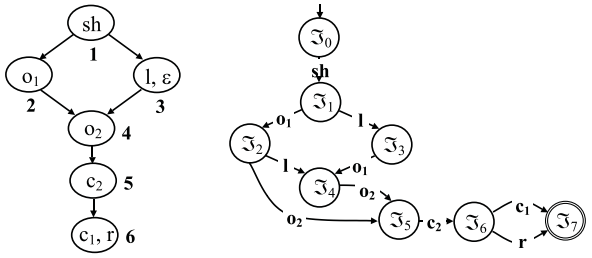

Figure 3: Observation graph and relevant index space.

Example 6. A ruler $\mathcal{R}_\psi$ for Ψ can be defined by the set of faulty transitions, namely $\mathcal{R}_\psi = \{(T_0(p), s), (T_3(b_1), \mathit{fo}_1), (T_4(b_1), \mathit{fc}_1), (T_3(b_2), \mathit{fo}_2), (T_4(b_2, \mathit{fc}_2)\}$. □

A *diagnostic problem* is a 4-tuple $\wp(\Sigma) = (\Sigma_0, \mathcal{V}, \mathcal{O}, \mathcal{R})$, where Σ_0 is the initial state of Σ, $\mathcal{V}$ a viewer, $\mathcal{O}$ a fragmented observation of Σ, and $\mathcal{R}$ a ruler. A *sub-problem* $\wp(\Sigma)_{[i]}$, $i \in [0\,..\,n]$, is the diagnostic problem relevant to the sub-observation $\mathcal{O}_{[i]}$. A *candidate diagnosis* of $\wp(\Sigma)$ is a generic diagnosis $\delta = h \otimes \mathcal{R}$ where h is a history in $\mathit{Bhv}(\Sigma, \Sigma_0)$ such that $h \boxtimes \mathcal{V} \in \|\mathfrak{I}(\mathcal{O})\|$. The *static solution* of $\wp(\Sigma)$, $\Delta(\wp(\Sigma))$, is the set of candidate diagnoses of $\wp(\Sigma)$. The *dynamic solution* of $\wp(\Sigma)$, $\mathbf{\Delta}(\wp(\Sigma))$, is the sequence of static solutions relevant to all the sub-problems of $\wp(\Sigma)$:

$$\mathbf{\Delta}(\wp(\Sigma)) = \left\langle \Delta(\wp(\Sigma)_{[0]}), \Delta(\wp(\Sigma)_{[1]}), \ldots, \Delta(\wp(\Sigma)_{[n]}) \right\rangle .$$

Example 7. Consider $\wp(\Psi) = (\Psi_0, \mathcal{V}_\psi, \mathcal{O}_\psi, \mathcal{R}_\psi)$. Based on $\mathit{Bhv}(\Psi, \Psi_0)$ in Fig. 2 and $\mathfrak{I}(\mathcal{O}_\psi)$ in Fig. 3, it can be shown that $\Delta(\wp(\Psi)) = \{\{s\}\}$, which corresponds to the only history $h = \langle T_0(p), T_1(b_1), T_1(b_2), T_1(p), T_2(b_2), T_2(b_1) \rangle$: a short circuit has occurred on the line and the protection apparatus has reacted correctly. □

5 Problem Solving

During its operation, Σ reacts to external events and generates observable events that are perceived as a fragmented observation. Solving a diagnostic problem means generating its dynamic solution based on each new message. The challenge is to generate each $\Delta(\wp(\Sigma)_{[i]})$ incrementally, based on the previous solution $\Delta(\wp(\Sigma)_{[i-1]})$ and the new message μ_i, avoiding performing the entire model-based reasoning rooted in the initial state. Thus, we need to know the state reached by the system at the occurrence of each message and such a knowledge has to be drawn directly from the component models. How to generate, by means of a divide-and-conquer algorithm, the evolution of a system, given a state from which to start from and an observed label, was dealt with in previous works by the authors [Lamperti and Zanella, 2003]. But, even in case the previous state were known, the current state is bound to be uncertain owing to silent transitions and the uncertainty of messages. However, the set of possible states at each newly generated message is confined within a limited domain, this corresponding to all the states reachable via silent transitions.

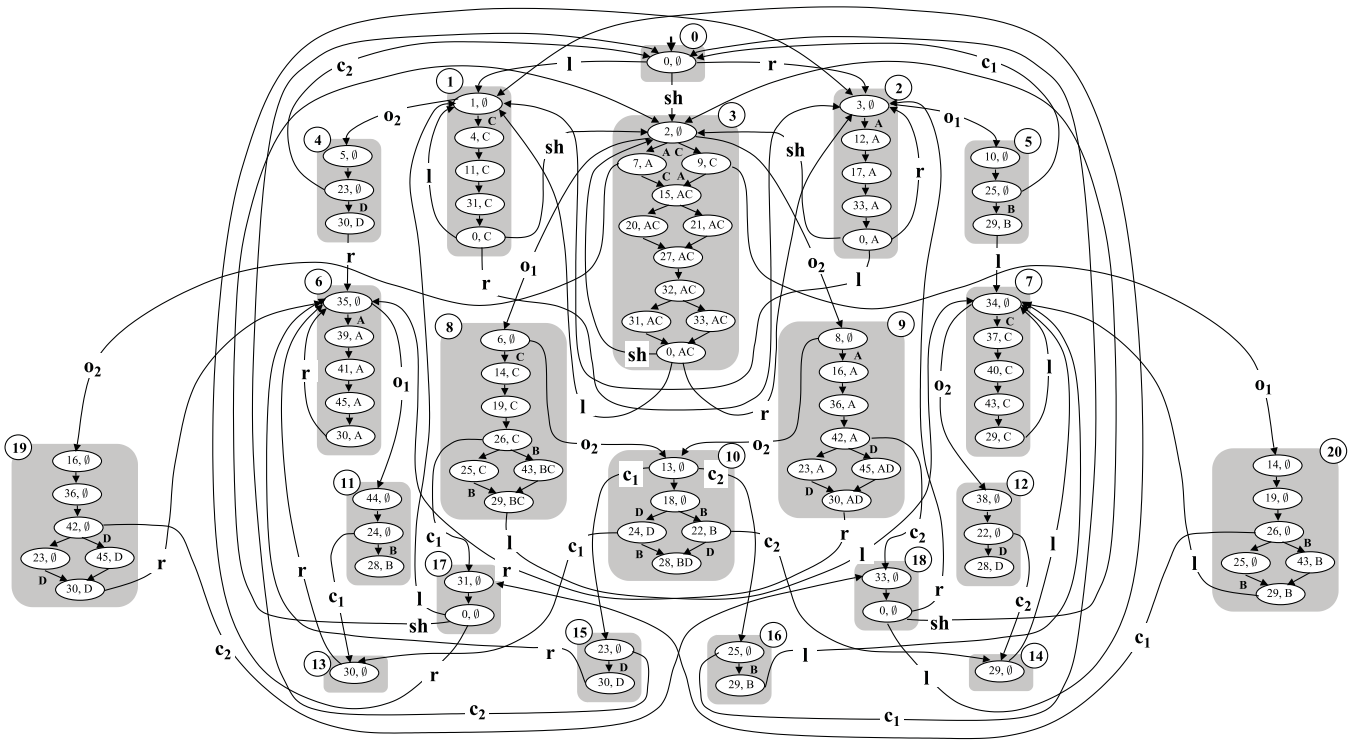

Figure 4: Monitor $Mtr(\Psi, \Psi_0, \mathcal{V}_\psi, \mathcal{R}_\psi)$.

Let σ_0 be a node of $Bhv(\Sigma, \Sigma_0)$, $\mathcal{V}$ a viewer, and $\mathcal{R}$ a ruler for Σ. The *diagnostic closure* $Dcl(\sigma_0, \mathcal{V}, \mathcal{R}) = (\mathbf{S}, \mathbf{E}, \mathbf{T}, S_0, \mathbf{S}_{\text{out}})$ is an automaton such that $S_0 = (\sigma_0, \mathfrak{D}_0)$ is the *root*, and each state $S \in \mathbf{S}$ is a pair $(\sigma, \mathfrak{D})$ where σ is a state of $Bhv(\Sigma, \Sigma_0)$ and $\mathfrak{D}$ the *candidate attribute*, namely a set of diagnoses $\delta = h \otimes \mathcal{R}$ where $h = \sigma_0 \rightsquigarrow \sigma$ is a history segment in $Bhv(\Sigma, \Sigma_0)$. $\mathbf{E}$ is the set of transitions of Σ. $\mathbf{T} : \mathbf{S} \times \mathbf{E} \mapsto \mathbf{S}$ is the transition function such that $(\sigma, \mathfrak{D}) \xrightarrow{T} (\sigma', \mathfrak{D}') \in \mathbf{T}$ iff T is a silent transition of Σ in $\mathcal{V}$ and $\sigma \xrightarrow{T} \sigma'$ is a transition in $Bhv(\Sigma, \Sigma_0)$. $\mathbf{S}_{\text{out}} \subseteq \mathbf{S}$ is the *leaving set*, where $S = (\sigma, \mathfrak{D}) \in \mathbf{S}_{\text{out}}$ iff there exists a transition $\sigma \xrightarrow{T'} \sigma'$ in $Bhv(\Sigma, \Sigma_0)$ where T' is visible in $\mathcal{V}$.

Example 8. With reference to $Bhv(\Psi, \Psi_0)$ in Fig. 2, $Dcl(5, \mathcal{V}_\psi, \mathcal{R}_\psi)$ is the subgraph involving states $(5, \emptyset)$, $(23, \emptyset)$, and $(30, \{\{fc_2\}\})$, with $\mathbf{S}_{\text{out}} = \{(23, \emptyset), (30, \{\{fc_2\}\})\}$, whose states in the behavior are left by visible transitions $T_2(b_2)$ and $T_3(p)$, respectively. □

The *monitor* of a system Σ with initial state Σ_0, a viewer $\mathcal{V}$, and a ruler $\mathcal{R}$ is a graph $Mtr(\Sigma, \Sigma_0, \mathcal{V}, \mathcal{R}) = (\mathcal{N}, \mathcal{L}, \mathcal{E}, N_0)$ where $\mathcal{N}$ is the set of nodes, $\mathcal{L}$ the set of labels, $\mathcal{E}$ the set of edges, and N_0 the initial node. Each node $N \in \mathcal{N}$ is the diagnostic closure of a state $S_0 \in Bhv(\Sigma, \Sigma_0)$, $N = Dcl(S_0, \mathcal{V}, \mathcal{R}) = (\mathbf{S}, \mathbf{E}, \mathbf{T}, S_0, \mathbf{S}_{\text{out}})$. Let $\mathbb{S}_{\text{out}} = \bigcup_{N \in \mathcal{N}} \mathbf{S}_{\text{out}}(N)$, $\mathbb{S}_0 = \bigcup_{N \in \mathcal{N}} \{S_0(N)\}$, and $\mathbf{V}$ and $\mathbf{R}$ the domains of labels in $\mathcal{V}$ and $\mathcal{R}$, respectively. Each edge $E \in \mathcal{E}$ is marked by a label in $\mathbb{S}_{\text{out}} \times (\mathbf{V} - \{\varepsilon\}) \times \mathbf{R} \times \mathbb{S}_0$. An edge $N \xrightarrow{(S, \ell, \varphi, S')} N'$, where $S = (\sigma, \mathfrak{D})$ and $S' = (\sigma', \mathfrak{D}')$ are internal nodes of N and N', respectively, is such that (i) S' is the root of N', (ii) $\sigma \xrightarrow{T} \sigma'$ is a transition in $Bhv(\Sigma, \Sigma_0)$, (iii) ℓ is the (visible) label associated with T in $\mathcal{V}$, and (iv) φ is the label associated with T in $\mathcal{R}$. The initial node N_0 is such that $S_0(N_0) = (\Sigma_0, \mathfrak{D}_0)$. Let N be a node of $Mtr(\Sigma, \Sigma_0, \mathcal{V}, \mathcal{R})$. The *local candidate set* $\Delta^{\text{loc}}(N)$ is the union of the candidate attributes relevant to the internal states of N.

Example 9. Shown in Fig. 4 is $Mtr(\Psi, \Psi_0, \mathcal{V}_\psi, \mathcal{R}_\psi)$. Each node of the monitor is confined by a shaded box and labeled by $i \in [0\,..\,20]$ (standing for N_i), where 0 is the root. Within each node, faulty transitions are marked by letters A, B, C, or D, which are a shorthand for faults *fo*$_1$, *fc*$_1$, *fo*$_2$, and *fc*$_2$, respectively. Candidate attributes are written as strings of such letters, e.g., AC is a shorthand for $\{\{A, C\}\}$. Edges between nodes are arrows from an internal state of the leaving node to the root of the entering node, and marked by the label in $\mathcal{V}_\psi$. Identifiers of component transitions are omitted (see Fig. 2). $T_0(p)$ is the only transition both visible (label *sh*) and faulty (label *s*). Its ruler label is omitted. □

The notion of a monitor allows the tracing of the system states based on a given fragmented observation. However, such a state is uncertain for three reasons: (i) The uncertain nature of the message, (ii) The unobservability of the transitions within the nodes of the monitor, and (iii) The nondeterminism of the monitor, where different edges leaving the same node can be marked by the same label.

The *diagnostic join* of two non-empty diagnostic sets Δ_1 and Δ_2 is the diagnostic set $\Delta_1 \bowtie \Delta_2 = \{\delta \mid \delta = \delta_1 \cup \delta_2, \delta_1 \in \Delta_1, \delta_2 \in \Delta_2\}$. The *diagnostic union* of a non-empty diagnostic set Δ and a label $\varphi \in \mathbf{R}$, is the diagnostic set

$$\Delta \uplus \varphi = \begin{cases} \Delta & \text{if } \varphi = \varepsilon \\ \{\delta' \mid \delta' = \delta \cup \{\varphi\}, \delta \in \Delta\} & \text{otherwise.} \end{cases}$$

Let $\wp(\Sigma) = (\Sigma_0, \mathcal{V}, \mathcal{O}, \mathcal{R})$, where $\mathcal{O} = \langle \ell_1, \ldots, \ell_n \rangle$ is

plain, and $Mtr(\Sigma, \Sigma_0, \mathcal{V}, \mathcal{R}) = (\mathcal{N}, \mathcal{L}, \mathcal{E}, N_0)$. A *context* $\chi = (N, \Delta)$ is an association between a node $N \in \mathcal{N}$ and a diagnostic set Δ. A *monitoring state* $\mathcal{M}$ is a set of contexts. The *trajectory* of $\wp(\Sigma)$, $Trj(\wp(\Sigma))$, is a sequence $\langle \mathcal{M}_0, \mathcal{M}_1, \ldots, \mathcal{M}_n \rangle$ of monitoring states defined as follows: $\mathcal{M}_0 = \{(N_0, \{\emptyset\})\}$; $\forall i \in [1\,..\,n]$, $\mathcal{M}_i$ is the minimal set of contexts $\chi' = (N', \Delta')$ such that $\chi \in \mathcal{M}_{i-1}$, $\chi = (N, \Delta)$, $S \in \mathbf{S}_{\text{out}}(N)$, $S = (\sigma, \mathfrak{D})$, $N \xrightarrow{(S, \ell_i, \varphi, S_0(N'))} N' \in \mathcal{L}$, and $\Delta' \supseteq (\mathfrak{D} \bowtie \Delta) \uplus \varphi$.

Example 10. Considering $Mtr(\Psi, \Psi_0, \mathcal{V}_\psi, \mathcal{R}_\psi)$ in Fig. 4, assume $\mathcal{O}'_\psi = \langle sh, o_1, l \rangle$ and $\wp'(\Psi) = (\Psi_0, \mathcal{V}_\psi, \mathcal{O}'_\psi, \mathcal{R}_\psi)$. Then, $Trj(\wp'(\Psi)) = \langle \mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3 \rangle$, where $\mathcal{M}_0 = \{(N_0, \{\emptyset\})\}$, $\mathcal{M}_1 = \{(N_3, \{\{s\}\})\}$, $\mathcal{M}_2 = \{(N_8, \{\{s\}\}), (N_{20}, \{\{s, C\}\})\}$, and $\mathcal{M}_3 = \{(N_7, \{\{s, B, C\}\})\}$. □

The *candidate sequence* of $\wp(\Sigma)$ is a list of diagnostic sets, $Cand(\wp(\Sigma)) = \langle \Delta_0, \Delta_1, \ldots, \Delta_n \rangle$, where

$$\forall i \in [0\,..\,n] \left(\Delta_i = \bigcup_{(N, \Delta) \in \mathcal{M}_i} (\Delta^{\text{loc}}(N) \bowtie \Delta) \right).$$

Example 11. Considering Example 10, $Cand(\wp'(\Psi)) = \langle \Delta_0, \Delta_1, \Delta_2, \Delta_3 \rangle$, where $\Delta_0 = \{\emptyset\}$, $\Delta_1 = \{\{s\}, \{s, A\}, \{s, C\}, \{s, A, C\}\})$, $\Delta_2 = \{\{s\}, \{s, C\}, \{s, B, C\}\}$, $\Delta_3 = \{\{s, B, C\}\}$. □

The notions of trajectory and candidate sequence were introduced based on plain observations. On the other hand, monitoring-based diagnosis is meant for diagnostic problems with fragmented observations. Such observations are represented by a DAG from which an index space can be generated. Since each state of the index space corresponds to several possible ways in which observable labels may have been generated by the evolution of Σ (several plain observations), the computation of the candidate sequence in the general case requires associating each state $\Im$ of the index space with the set of monitoring states that are consistent with all the plain observations relevant to $\Im$.

Let $\wp(\Sigma) = (\Sigma_0, \mathcal{V}, \mathcal{O}, \mathcal{R})$, $\mathfrak{I}(\mathcal{O}) = (\mathbb{S}, \mathbb{E}, \mathbb{T}, S_0, \mathbb{S}_{\text{f}})$. The *decoration* of $\mathfrak{I}(\mathcal{O})$ based on $\wp(\Sigma)$ is an automaton $\mathfrak{I}^{\mathcal{M}}(\mathcal{O}) = (\mathbb{S}^{\mathcal{M}}, \mathbb{E}^{\mathcal{M}}, \mathbb{T}^{\mathcal{M}}, S_0^{\mathcal{M}}, \mathbb{S}_{\text{f}}^{\mathcal{M}})$ isomorphic to $\mathfrak{I}(\mathcal{O})$, where each state $\Im \in \mathbb{S}$ is marked by a *monitoring attribute* $\mathcal{M} = \bigcup_{\mathcal{O}' \in \|\Im\|} \mathcal{M}_k$, where $\|\Im\|$ is the set of plain observations up to $\Im$ in $\mathfrak{I}(\mathcal{O})$, $\mathcal{O}' = \langle \ell_1, \ldots, \ell_k \rangle$, $\wp'(\Sigma) = (\Sigma_0, \mathcal{V}, \mathcal{O}', \mathcal{R})$, and $Trj(\wp'(\Sigma)) = \langle \mathcal{M}_0, \mathcal{M}_1, \ldots, \mathcal{M}_k \rangle$.

Example 12. Consider $\wp(\Psi) = (\Psi_0, \mathcal{V}_\psi, \mathcal{O}_\psi, \mathcal{R}_\psi)$ (see Fig. 3). The decoration of $\mathfrak{I}(\mathcal{O}_\psi)$ can be expressed by determining each monitoring attribute $\mathcal{M}_i$ that is relevant to node $\Im_i$, $i \in [0\,..\,7]$: $\mathcal{M}_0 = \{(N_0, \{\emptyset\})\}$, $\mathcal{M}_1 = \{(N_3, \{\{s\}\})\}$, $\mathcal{M}_2 = \{(N_8, \{\{s\}\}), (N_{20}, \{\{s, C\}\})\}$, $\mathcal{M}_3 = \{(N_1, \{\{s, A, C\}\})\}$, $\mathcal{M}_4 = \{(N_7, \{\{s, B, C\}\})\}$, $\mathcal{M}_5 = \{(N_{10}, \{\{s\}\}), (N_{12}, \{\{s, B, C\}\})\}$, $\mathcal{M}_6 = \{(N_{16}, \{\{s\}\}), (N_{14}, \{\{s, B\}, \{s, B, C\}\})\}$, $\mathcal{M}_7 = \{(N_0, \{\{s\}\})\}$. □

Based on the concept of index-space decoration, both notions of trajectory and candidate sequence can be straightforwardly generalized to diagnostic problems involving a fragmented observation with uncertain messages as follows. Let $\wp(\Sigma)$ be a diagnostic problem involving a fragmented observation $\mathcal{O} = \langle \mu_1, \ldots, \mu_n \rangle$. Let $\mathfrak{I}^{\mathcal{M}}(\mathcal{O}_{[i]}) = (\mathbb{S}_i^{\mathcal{M}}, \mathbb{E}_i^{\mathcal{M}}, \mathbb{T}_i^{\mathcal{M}}, S_{0_i}^{\mathcal{M}}, \mathbb{S}_{\text{f}_i}^{\mathcal{M}})$, $i \in [0\,..\,n]$. The (*generalized*) *trajectory* of $\wp(\Sigma)$ is the sequence of monitoring states

i	$\mathbb{S}_{\text{f}_i}^{\mathcal{M}}$	Δ_i
0	$\{\Im_0\}$	$\{\emptyset\}$
1	$\{\Im_1\}$	$\{\{s\}, \{s, A\}, \{s, C\}, \{s, A, C\}\}$
2	$\{\Im_2\}$	$\{\{s\}, \{s, C\}, \{s, B, C\}\}$
3	$\{\Im_2, \Im_4\}$	$\{\{s\}, \{s, C\}, \{s, B, C\}\}$
4	$\{\Im_5\}$	$\{\{s\}, \{s, B\}, \{s, D\}, \{s, B, C\}, \{s, B, D\}, \{s, B, C, D\}\}$
5	$\{\Im_6\}$	$\{\{s\}, \{s, B\}, \{s, B, C\}\}$
6	$\{\Im_7\}$	$\{\{s\}\}$

Table 1: Generation of the candidate sequence $\mathbf{\Delta}(\wp(\Psi))$.

$Trj(\wp(\Sigma)) = \langle \mathcal{M}_0, \mathcal{M}_1, \ldots, \mathcal{M}_n \rangle$, where

$$\forall i \in [0\,..\,n] \left(\mathcal{M}_i = \bigcup_{(\Im, \mathcal{M}) \in \mathbb{S}_{\text{f}_i}^{\mathcal{M}}} \mathcal{M} \right).$$

The definition of the candidate sequence does not change, as each Δ_i depends on the monitoring state $\mathcal{M}_i$ in the trajectory.

Theorem 1. *The candidate sequence is the dynamic solution of the diagnostic problem:* $Cand(\wp(\Sigma)) = \mathbf{\Delta}(\wp(\Sigma))$.

Theorem 1 is the formal foundation of the monitoring-based diagnostic technique. The static solution generated at each processing step consists in a sound and complete set of candidate diagnoses with respect to the (uncertain) messages received so far. The proof is omitted for space reasons.

Example 13. With reference to the diagnostic problem $\wp(\Psi)$ defined in Example 12, the candidate sequence $\mathbf{\Delta}(\wp(\Psi))$ will be $\langle \Delta_0, \Delta_1, \ldots, \Delta_6 \rangle$, as detailed in Table 1. Specifically, each sub-observation $\mathcal{O}_{[i]}$ is associated with the set of final states $\mathbb{S}_{\text{f}_i}^{\mathcal{M}}$ of the decoration $\mathfrak{I}^{\mathcal{M}}(\mathcal{O}_{[i]})$, whose monitoring attributes were computed in Example 12. The diagnostic set reduces to the singleton $\{\{s\}\}$ upon the arrival of the sixth message: although both N_{14} and N_{16} are relevant to the monitoring attribute $\mathcal{M}_6$, attribute $\mathcal{M}_7$ does not include any node of the monitor leaving N_{14}, but only N_0, a neighbor of N_{16}. As expected, Δ_6 equals the static solution $\Delta(\wp(\Psi))$ obtained in Example 7. More generally, and in accordance with Theorem 1, it can be shown that $Cand(\wp(\Psi)) = \mathbf{\Delta}(\wp(\Psi))$. □

The diagnostic technique has been substantiated by a variety of algorithms. It is worth mentioning the *Increment* procedure, which builds the index space. The peculiarity of *Increment* is twofold: (1) the new index space is generated incrementally, upon the reception of each message and based on the previous index space, and (2) such a generation is performed directly, without any transformation from a nondeterministic to a deterministic automaton.

6 Conclusion

This paper deals with monitoring-based diagnosis of DESs, a task that is also considered by the diagnoser approach [Sampath *et al.*, 1995; 1996] and its extension [Rozé and Cordier, 2002], by the incremental decentralized diagnoser approach [Pencolé *et al.*, 2001], and by the bridged diagnostic method [Lamperti and Zanella, 2004a]. All these contributions differ from the current one in several aspects.

First, in the class of considered systems: [Sampath *et al.*, 1995; 1996; Pencolé *et al.*, 2001] deal with synchronous DESs, [Rozé and Cordier, 2002] with timed asynchronous

DESs (and is oriented to telecommunication networks), while [Lamperti and Zanella, 2004a] with polymorphic systems, integrating both synchronous and asynchronous behavior. The method in this paper, instead, copes with untimed asynchronous DESs, where a system transition amounts to a component transition, triggered by a single event, and wherein there may be delays between the time an event is received and the time it is consumed. Every such system may follow behavioral silent cycles over time, which is not the case for the diagnoser approach and its extension.

Second, monitoring-based diagnosis in this paper adapts the diagnostic algorithm in [Lamperti and Zanella, 2004a], which is quite different from the algorithms of both the diagnoser approach and the incremental decentralized diagnoser approach. Moreover, the definition of a diagnostic problem in the current paper differs from [Lamperti and Zanella, 2004a] since it includes the notions of a ruler and a viewer [Lamperti and Zanella, 2004b], which decouple the component models from the descriptions of their observability and abnormality properties.

The essential novelty of this paper lies in the extension of monitoring-based diagnosis to uncertain observations. In [Rozé and Cordier, 2002] the observation is a completely certain stream of time-stamped alarms and represents time constraints explicitly. A certain plain observation is considered in [Sampath *et al.*, 1995; 1996; Pencolé *et al.*, 2001; Lamperti and Zanella, 2004a]. The approach in this paper, instead, takes as input a sequence of observation fragments, each fragment being both logically and temporally uncertain. At the occurrence of each new fragment, the index space is updated and a new hyperstate of the monitor is generated. Owing to temporal uncertainty, at every newly-received message, additional sequences of labels may have to be added to the ones hypothesized so far. However, the assumption inherent to fragmented observations prevents any sequence of labels hypothesized in previous monitoring steps from being refuted. The algorithm for updating the index space, which is not shown in this paper for shortness, produces as output a deterministic automaton in one shot. In [Lamperti and Zanella, 2002], instead, the index space was built based on the whole uncertain observation of the a posteriori diagnosis problem and its construction required the transformation of a nondeterministic automaton into a deterministic one.

Each new hyperstate of the monitor can be built modularly by breaking down the problem into a hierarchy of independent subproblems, where parallelism may be exploited and no previous construction of the global system behavior is needed, the same as in previous contributions by the authors [Lamperti and Zanella, 2003].

A major limitation of the approach is its computational complexity. The tractability of state estimation was addressed by Livingstone [Muscettola *et al.*, 1998] and L2 [Kurien and Nayak, 2000] by adopting optimizations and approximations. In this paper, instead, an exact method is applied and the scalability of the approach is increased by distributed and incremental processing, in the creation of both the monitor and the index space. Moreover, the efficiency of the reasoning mechanism could benefit from a trade-off between time and space, in particular if some nodes of the monitor (or some compiled knowledge supporting their construction) were generated offline by preprocessing utilities. Providing experimental evidence to these claims is an engagement for future work.

References

[Brusoni *et al.*, 1998] 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.

[Console *et al.*, 2002] L. Console, C. Picardi, and M. Ribaudo. Process algebras for systems diagnosis. *Artificial Intelligence*, 142(1):19–51, 2002.

[Grastien *et al.*, 2004] A. Grastien, M.O. Cordier, and C. Largouët. Extending decentralized discrete-event modelling to diagnose reconfigurable systems. In *Fifteenth International Workshop on Principles of Diagnosis –DX'04*, pages 75–80, Carcassonne, F, 2004.

[Kurien and Nayak, 2000] J. Kurien and P.P. Nayak. Back to the future for consistency-based trajectory tracking. In *Eleventh International Workshop on Principles of Diagnosis – DX'00* , pages 92–100, Morelia, MX, 2000.

[Lamperti and Zanella, 2002] G. Lamperti and M. Zanella. Diagnosis of discrete-event systems from uncertain temporal observations. *Artificial Intelligence*, 137(1–2):91–163, 2002.

[Lamperti and Zanella, 2003] G. Lamperti and M. Zanella. *Diagnosis of Active Systems – Principles and Techniques* , volume 741 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publisher, Dordrecht, NL, 2003.

[Lamperti and Zanella, 2004a] G. Lamperti and M. Zanella. A bridged diagnostic method for the monitoring of polymorphic discrete-event systems. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics* , 34(5):2222–2244, 2004.

[Lamperti and Zanella, 2004b] G. Lamperti and M. Zanella. Diagnosis of discrete-event systems by separation of concerns, knowledge compilation, and reuse. In *Sixteenth European Conference on Artificial Intelligence – ECAI'2004*, pages 838–842, Valencia, E, 2004.

[Muscettola *et al.*, 1998] N. Muscettola, P.P. Nayak, B. Pell, and B.C. Williams. Remote Agent: to boldly go where no AI system has gone before. *Artificial Intelligence*, 103(1–2):5–47, 1998.

[Pencolé *et al.*, 2001] Y. Pencolé, M.O. Cordier, and L. Rozé. Incremental decentralized diagnosis approach for the supervision of a telecommunication network. In *Twelfth International Workshop on Principles of Diagnosis – DX'01* , pages 151–158, San Sicario, I, 2001.

[Pencolé, 2004] Y. Pencolé. Diagnosability analysis of distributed discrete event systems. In *Sixteenth European Conference on Artificial Intelligence –ECAI'2004*, pages 43–47, Valencia, E, 2004.

[Rozé and Cordier, 2002] L. Rozé and M.O. Cordier. Diagnosing discrete-event systems: extending the 'diagnoser approach 'to deal with telecommunication networks. *Journal of Discrete Event Dynamic Systems: Theory and Application* , 12:43–81, 2002.

[Sampath *et al.*, 1995] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D.C. Teneketzis. Diagnosability of discrete-event systems. *IEEE Transactions on Automatic Control*, 40(9):1555–1575, 1995.

[Sampath *et al.*, 1996] 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.

[Schumann *et al.*, 2004] A. Schumann, Y. Pencolé, and S. Thiébaux. Diagnosis of discrete-event systems using binary decision diagrams. In *Fifteenth International Workshop on Principles of Diagnosis –DX'04*, pages 197–202, Carcassonne, F, 2004.

[Struss, 1997] P. Struss. Fundamentals of model-based diagnosis of dynamic systems. In *Fifteenth International Joint Conference on Artificial Intelligence – IJCAI'97*, pages 480–485, Nagoya, J, 1997.