Automatica 49 (2013) 223–231

Contents lists available at SciVerse ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Robust diagnosis of discrete-event systems against permanent loss of observations☆

Lilian K. Carvalho [a], Marcos V. Moreira [a], João C. Basilio [a,1], Stéphane Lafortune [b]

[a] Universidade Federal do Rio de Janeiro, COPPE - Programa de Engenharia Elétrica, 21949-900, Rio de Janeiro, RJ, Brazil

[b] Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:
Received 20 July 2011
Received in revised form
20 April 2012
Accepted 30 July 2012
Available online 9 October 2012

Keywords:
Discrete-event systems
Fault diagnosis
Sensor failures
Robust diagnosability

ABSTRACT

We consider the problem of diagnosing the occurrence of a certain unobservable event of interest, the *fault* event, in the operation of a partially-observed discrete-event system subject to permanent loss of observations modeled by a finite-state automaton. Specifically, it is assumed that certain sensors for events that would *a priori* be observable may fail at the outset, thereby resulting in a loss of observable events; the diagnostic engine is not directly aware of such sensor failures. We explore a previous definition of robust diagnosability of a given fault event despite the possibility of permanent (and unknown *a priori*) loss of observations and present a polynomial time verification algorithm to verify robust diagnosability and a methodology to perform online diagnosis in this scenario using a set of partial diagnosers.

1. Introduction

The basic event diagnosis problem for discrete-event systems is to perform model-based inferencing at run-time, using sequences of observable events, and determine, with certainty, if a given unobservable "fault" event has occurred or not in the past. The property of diagnosability formally captures the ability to always detect at run-time any occurrence of the given fault event, within a finite number of event transitions. There is a very large body of literature on (offline) diagnosability analysis and (online) event diagnosis of discrete-event systems modeled by automata, the modeling formalism considered in this paper; see, e.g., Boel and van Schuppen (2002), Debouk, Lafortune, and Teneketzis (2000), Genc (2008), Jéron, Marchand, Pinchinat, and Cordier (2006), Kumar and Takai (2009), Lin (1994), Lunze and Schröder (2004), Pencolé and Cordier (2005), Qiu and Kumar (2006), Sampath, Sengupta, Lafortune, Sinnamohideen, and Teneketzis (1995), Thorsley and Teneketzis (2005), Tripakis (2002), Wang, Yoo, and Lafortune (2007), Ye, Dague, and Yan (2009), Yoo and Lafortune (2002), Zad, Kwong, and Wonham (2003) and the references contained therein. Two classes of automata derived from the automaton model of the system have been defined in the above works: diagnosers and verifiers. Both diagnosers and verifiers can be used for offline analysis of diagnosability properties; online diagnosis is usually implemented using diagnosers.

Let us assume that the given set of sensors attached to the system is recording all potentially observable events at run-time. We are interested in the situation where sensors for some combinations of (potentially observable) events fail prior to the first occurrence of an event they are monitoring; such failures are assumed to be permanent and unknown *a priori*. In this case, if online diagnosis is performed using a standard diagnoser built on the basis of all potentially observable events, then this diagnoser could get stuck in some states (e.g., no further observed event, or occurrence of an event not in the current active event set) or could even issue incorrect diagnostic decisions; an example is presented in Section 3. We would like to still perform correct diagnosis of the original unobservable fault event despite the (unknown *a priori*) loss of observations resulting from sensor failures.

Recently, there have been some works on sensor failures in supervisory control of discrete-event systems (see, e.g., Rohloff (2005); Sanchez and Montoya (2006)), on various notions of "robust" diagnosis of discrete-event systems in the presence of potentially faulty sensors, in particular, Basilio and Lafortune (2009), Carvalho, Basilio, and Moreira (2010, 2012), Contant, Lafortune, and Teneketzis (2006), and Takai (2010, 2012) and on

☆ This work was partially supported by the Brazilian Research Council (CNPq) grant 200820/2006-0 and by the US National Science Foundation grant EECS-0624821. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Jan Komenda, under the direction of Editor Ian R. Petersen.

E-mail addresses: lilian@coep.ufrj.br (L.K. Carvalho), moreira@dee.ufrj.br (M.V. Moreira), basilio@poli.ufrj.br (J.C. Basilio), stephane@umich.edu (S. Lafortune).

[1] Tel.: +55 21 2562 8021; fax: +55 21 2562 8627.

0005-1098/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2012.09.017

fault diagnosis under unreliable observations (Athanasopoulou, Lingxi, & Hadjicostis, 2010; Thorsley, Yoo, & Garcia, 2008).

In this paper, we deal with the problem of robust diagnosis against permanent loss of observations. This problem was first introduced by Lima, Basilio, Lafortune, and Moreira (2010), and can be stated as follows. Let us assume that a given unobservable fault event, σ_f, is diagnosable in a given system for the set of all observable events Σ_o, in the sense of Sampath et al. (1995). Let $\Sigma'_o \subset \Sigma_o$ be a proper subset of Σ_o for which diagnosability still holds. Then Σ'_o is called a diagnosis basis (Basilio, Lima, Lafortune, & Moreira, 2012) and the events in the set $\Sigma_o \setminus \Sigma'_o$ are said to be *redundant*; we call $\Sigma'_{uo} := \Sigma_o \setminus \Sigma'_o$ the set of redundant events associated with Σ'_o; the partial diagnoser built for Σ'_o does not record these (potentially observable) events. Lima et al. (2010) present a necessary and sufficient condition for robust diagnosability against permanent sensor failures using a union diagnoser, i.e., a diagnoser that accepts the union of the languages of all partial diagnosers formed with all sets Σ'_o that are diagnosis bases. It is not difficult to see that union diagnosers tend to have huge state spaces, which makes the verification test very computationally demanding.

In order to overcome the potential state space explosion of union diagnosers, we propose in this paper an offline test based on the use of a special type of verifier automata. This procedure avoids the worst-case exponential complexity of diagnosers, as verifiers can be computed in the worst-case polynomial time in the size of the system. We also discuss how to perform online diagnosis for systems that are robust diagnosable.

This paper is structured as follows. In Section 2, we present some background on fault diagnosis of discrete-event systems. In Section 3, we present the definition of robust diagnosability against permanent loss of observations. In Section 4 we develop an offline test for the verification of robust diagnosability and in Section 5 we discuss the online implementation of robust diagnosers. Finally, conclusions are drawn in Section 6.

2. Preliminaries

Let

$$G = (X, \Sigma, f, \Gamma, x_0), \tag{1}$$

be a deterministic automaton, where X denotes the state space, Σ the event set, $f : X \times \Sigma \to X$ the state transition function, which is partially defined over its domain, Γ the active event set, and x_0 the initial state. Let us partition Σ as $\Sigma = \Sigma_o \dot{\cup} \Sigma_{uo}$, i.e., $\Sigma = \Sigma_o \cup \Sigma_{uo}$, $\Sigma_o \cap \Sigma_{uo} = \emptyset$ and $\Sigma_{uo} \neq \emptyset$, where Σ_o and Σ_{uo} are, respectively, the set of observable and unobservable events, and let $\Sigma_f = \{\sigma_f\} \subseteq \Sigma_{uo}$ be a set whose unique element σ_f is the fault event to be detected. Finally, let us denote the language generated by G as L. We make the following common assumptions:

A1. Language L is live, i.e., $\Gamma(x_i) \neq \emptyset$ for all $x_i \in X$.

A2. There is no cycle of unobservable events in G.

The language L is said to be diagnosable if the occurrence of σ_f can be detected within a finite number of transitions after the occurrence of σ_f using only traces formed with events in Σ_o. Let the function $P_o : \Sigma^* \to \Sigma_o^*$ denote the standard natural projection that erases unobservable events; see Cassandras and Lafortune (2008). In addition, let P_o^{-1} denote the inverse projection of P_o, and assume that $L/s = \{t \in \Sigma^* : st \in L\}$ and that $\Psi(\Sigma_f)$ denotes the set of all traces of L that end with event σ_f. With some abuse of notation $\Sigma_f \in s$ denotes that $\bar{s} \cap \Psi(\Sigma_f) \neq \emptyset$. Language diagnosability can then be formally defined as follows (Sampath et al., 1995).

Definition 1. L is diagnosable with respect to $P_o : \Sigma^* \to \Sigma_o^*$ if, and only if, the following condition holds true:

$$(\exists n \in \mathbb{N})(\forall s \in \Psi(\Sigma_f))(\forall t \in L/s)(\|t\| \geq n \Rightarrow D),$$

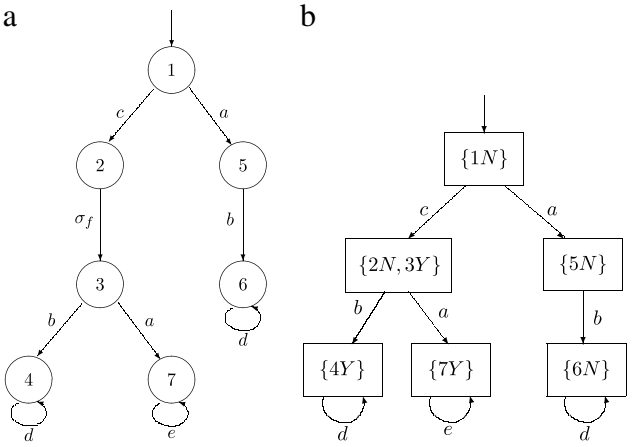

Fig. 1. Automaton G (a) and its diagnoser G_d (b).

where the diagnosability condition D is given as

$$(\forall \omega \in (P_o^{-1}(P_o(st)) \cap L))(\Sigma_f \in \omega).$$

3. Robust diagnosability against permanent loss of observations

Language diagnosability is usually performed in practice using diagnosers. A diagnoser is a deterministic automaton, which is built from the automaton that generates the language to be diagnosed and whose event set is formed with the observable events of G, and whose states are sets of states of G augmented by adding labels Y or N to indicate whether the fault event σ_f has occurred or not in reaching the state. In this regard, a state x_d of the diagnoser is called certain (or faulty) if $\ell = Y$ for all $x\ell \in x_d$, and normal (or non-faulty) if $\ell = N$ for all $x\ell \in x_d$. If there exist $x\ell, y\tilde{\ell} \in x_d$, x not necessarily distinct from y such that $\ell = Y$ and $\tilde{\ell} = N$, then x_d is an uncertain state of G_d. When the diagnoser reaches a certain (resp. normal) state, we are certain that the fault has occurred (resp. not occurred). However, when the diagnoser is in an uncertain state, we cannot draw any conclusion regarding the fault occurrence. If the diagnoser remains indefinitely in a cycle formed with uncertain states only, then it will not be possible to diagnose the fault occurrence.[2]

Fig. 1(a) shows the state transition diagram of an automaton G, for which $\Sigma = \{a, b, c, d, e, \sigma_f\}$, $\Sigma_o = \{a, b, c, d, e\}$, and $\Sigma_f = \{\sigma_f\}$. The corresponding diagnoser is depicted in Fig. 1(b). Notice that, since G_d has cycles in certain and normal states only, then we may say that L is diagnosable with respect to P_o and Σ_f. Indeed, if trace $s_Y = c\sigma_f bd^n$ ($n \in \mathbb{N}$) occurs, then the diagnoser goes from the initial state $\{1N\}$ to state $\{4Y\}$, indicating that the fault event σ_f has occurred. Assume now that a permanent loss of observation of the sensor that records the occurrence of event c took place before the first occurrence of c, and suppose that trace $s_Y = c\sigma_f ae^n, n \in \mathbb{N}$, has been generated. Since event σ_f is unobservable, the first event to be recognized by the diagnoser of Fig. 1(b) is a. When the diagnoser receives the information on the occurrence of a, it updates its state to $\{5N\}$, where it stands still since e is the only event that occurs next in trace s_Y and it is not in the active event set of $\{5N\}$. The diagnoser is, therefore, unable to process any further information it may receive regarding event occurrences, and so, it will not be able to reach a certain

[2] We refer the reader to Basilio et al. (2012) for a more detailed explanation about diagnosers and how they can be used as an offline test for diagnosability.