Automatica 61 (2015) 241–252

Contents lists available at ScienceDirect

Automatica

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

Codiagnosability and coobservability under dynamic observations: Transformation and verification☆

Xiang Yin [1], Stéphane Lafortune

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

ARTICLE INFO

Article history:
Received 9 December 2014
Received in revised form
30 April 2015
Accepted 7 August 2015
Available online 4 September 2015

Keywords:
Discrete-event systems
Fault diagnosis
Supervisory control
Diagnosability
Observability

ABSTRACT

We investigate the relationship between decentralized fault diagnosis and decentralized control of discrete event systems under dynamic observations. The key system-theoretic properties that arise in these problems are those of codiagnosability and coobservability, respectively. It was shown by Wang et al. (2011) that coobservability is transformable to codiagnosability; however, the transformation for the other direction has remained an open problem. In this paper, we consider a general language-based dynamic observations setting and show how the notion of K-codiagnosability can be transformed to coobservability. When the observation properties are transition-based, we present a new approach for the verification of transition-based codiagnosability. An upper bound of the diagnosis delay for decentralized diagnosis under transition-based observations is derived. Moreover, we show that transition-based [co]diagnosability is transformable to transition-based [co]observability. Our results thereby complement those in Wang et al. (2011) and provide a thorough characterization of the relationship between the two notions of codiagnosability and coobservability and their verification. In particular, our results allow the leveraging of the large existing literature on decentralized control synthesis to solve corresponding problems of decentralized fault diagnosis.

1. Introduction

Control and diagnosis are two important research areas in the study of Discrete Event Systems (DES). In complex automated systems, one is interested in designing a *supervisor* to restrict the system's behavior within a desired specification as well as designing a *diagnoser* in order to detect and isolate potential system faults. Due to limited sensing capabilities, both problems involve dealing with partial observation of the system's behavior. Moreover, many technological systems have decentralized information structures, thereby necessitating the development of decentralized control and diagnosis architectures, where a set of supervisors or diagnosers work as a team to ensure the desired specifications.

The property of *observability* arose in the study of the control of partially observed DES (Cieslak, Desclaux, Fawaz, & Varaiya, 1988; Lin & Wonham, 1988). It is well known that observability together with controllability provide the necessary and sufficient conditions for the existence of a supervisor that achieves a given specification. This notion was extended to *coobservability* for decentralized control problems, see, e.g., Overkamp and van Schuppen (2000), Rudie and Willems (1995), Rudie and Wonham (1992), Seatzu, Silva, and Van Schuppen (2013), Tripakis (2004) and Yoo and Lafortune (2002). Problems of centralized fault diagnosis of DES were initially studied in Lin (1994) and Sampath, Sengupta, Lafortune, Sinnamohideen, and Teneketzis (1995) where the notion of *diagnosability* was introduced and characterized. Several future investigations ensued and a large amount of literature has been published on this topic; the recent survey papers (Zaytoon & Lafortune, 2013; Zaytoon & Sayed-Mouchaweh, 2012) contain extensive bibliographies. Problems of decentralized fault diagnosis were initially considered in Debouk, Lafortune, and Teneketzis (2000), where several communication protocols were developed. In particular, in Protocol 3 of Debouk et al. (2000), all the local agents work independently, i.e., there is no communication among them. This protocol was further investigated in several subsequent works and the associated condition of *codiagnosability* was characterized and studied; see, e.g., Moreira, Jesus, and Basilio (2011), Qiu and Kumar (2006)

☆ This work was partially supported by the US National Science Foundation grants CCF-1138860 (Expeditions in Computing project ExCAPE: Expeditions in Computer Augmented Program Engineering), CNS-1446298, and CNS-1421122. The material in this paper was partially presented at the 2015 American Control Conference, July 1–3, 2015, Chicago, IL, USA. 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: xiangyin@umich.edu (X. Yin), stephane@umich.edu (S. Lafortune).

[1] Tel.: +1 7348343243; fax: +1 7347638041.

http://dx.doi.org/10.1016/j.automatica.2015.08.023
0005-1098/

and Wang, Yoo, and Lafortune (2007). State-based, distributed, and robust approaches to diagnosis have also been considered; see, e.g., Carvalho, Basilio, and Moreira (2012), Hashtrudi Zad, Kwong, and Wonham (2003), Pencolé and Cordier (2005), Seatzu et al. (2013), Su and Wonham (2005) and Zaytoon and Lafortune (2013).

All of the above-mentioned works are concerned with the case of *static* observations, where the set of observable events is fixed a priori. In many applications however, communication among different agents (see, e.g., Lin, 2014, Rudie, Lafortune, & Lin, 2003) as well as dynamic sensor activation (see, e.g., Cassez & Tripakis, 2008, Sears & Rudie, 2013a,b, Thorsley & Teneketzis, 2007, Wang, Lafortune, Girard, & Lin, 2010, Wang, Lafortune, Lin, & Girard, 2010) may lead to the case of *dynamic* observations. In the context of dynamic observations, the observability properties of an event are not fixed but may vary along each system trajectory. In Huang, Rudie, and Lin (2008), the authors studied the property of coobservability under dynamic observations. The fault diagnosis problem under dynamic observations has also been investigated in several works, such as Cassez and Tripakis (2008) and Thorsley and Teneketzis (2007) for the centralized case and Wang, Girard, Lafortune, and Lin (2011) for the decentralized case.

There is a wide literature on the two properties of coobservability and codiagnosability, due to their importance in solving decentralized control and diagnosis problems, respectively. However, almost all of the existing literature deals with problems of control and problems of diagnosis *separately*. An exception to this is the work in Wang et al. (2011), where it was shown, for the first time, how to map coobservability to codiagnosability, in the context of a language-based model for dynamic observations. This transformation from coobservability to codiagnosability makes it possible to leverage existing methodologies for solving (decentralized) diagnosis problems to solve (decentralized) control problems. However, to the best of our knowledge, the reverse transformation, from codiagnosability to coobservability, has remained an open problem, as mentioned in the recent survey (Sears & Rudie, 2015).

The contributions of this paper are two-fold. First, we show how to transform *K-codiagnosability to coobservability* under a general language-based dynamic observations setting. *K*-codiagnosability is a strong version of codiagnosability where it is required that any failure be diagnosed within *K* steps after its occurrence; in codiagnosability, the detection delay has to be finite but no *K* is specified. The transformation that we present exploits the fact that both the problem of *K*-codiagnosability and the problem of coobservability can be reduced to a *state disambiguation problem*. Second, we provide a new approach for the verification of transition-based codiagnosability. Our method is different from that in Wang et al. (2011) and adopts the standard verifier approach which is used for static diagnosis problem in the literature (Jiang, Huang, Chandra, & Kumar, 2001; Qiu & Kumar, 2006; Wang, Yoo et al., 2007; Yoo & Lafortune, 2002). Our approach ends up with the same complexity as the approach proposed in Wang et al. (2011); however it allows us to derive an upper bound for the maximal delay of diagnosis, which is not provided in Wang et al. (2011). Moreover, by using the derived upper bound for the maximal diagnosis delay, we show that transition-based [co]diagnosability is transformable to transition-based [co]observability. Therefore, the standard notion of diagnosability from Sampath et al. (1995) can be transformed to the standard notion of observability from Lin and Wonham (1988). Our results thereby complement those in Wang et al. (2011) and allow leveraging the large existing literature on problems of decentralized control to solve problems of decentralized fault diagnosis.

The remaining part of this paper is organized as follows. Section 2 presents necessary preliminaries and in particular it reviews the notions of codiagnosability and coobservability. In Section 3, the transformation from *K*-codiagnosability to coobservability under language-based observations is presented. In Section 4, we present a new approach for the verification of transition-based codiagnosability, with which an upper bound of the diagnosis delay for decentralized diagnosis under transition-based observations is derived. We illustrate the application of the transformation algorithm of Section 3 to sensor activation problems in Section 5. Finally, we conclude the paper in Section 6. Preliminary and partial versions of some of the results in Sections 3 and 5 are presented in Yin and Lafortune (2015).

2. Preliminaries

2.1. System model

We assume basic knowledge of DES and common notations (see, e.g., Cassandras & Lafortune, 2008). A DES is modeled as a deterministic finite-state automaton $G = (X^G, E^G, \delta^G, x_0^G)$, where X^G is the finite set of states, E^G is the finite set of events, $\delta^G : X^G \times E^G \to X^G$ is the partial transition function where $\delta^G(x, e) = y$ means that there is a transition labeled by event e from state x to state y, and $x_0^G \in X^G$ is the initial state. Function δ^G is extended to $X^G \times E^{G*}$ in the usual way. The behavior generated by G is described by $\mathcal{L}(G) = \{s \in E^{G*} : \delta^G(x_0^G, s)!\}$, where ! means "is defined". The set of transitions $TR(G)$ of G is defined by $TR(G) := \{(x, e) \in X^G \times E^G : \delta^G(x, e)!\}$. The prefix-closure of a language L is $\overline{L} = \{s \in E^{G*} : (\exists t \in E^{G*})[st \in L]\}$. We use notation $|\cdot|$ to denote the length of a string.

In both control and diagnosis problems, there are some local agents monitoring the plant based on their own observations. Here, we assume that there are n local agents and we denote by $\mathcal{I} = \{1, \ldots, n\}$ the index set of the local agents. In most of the existing literature, the observation properties of events are specified by natural projection operations, i.e., for each agent $i \in \mathcal{I}$, the set of observable events $E_{o,i} \subset E^G$ is fixed a priori. We denote by $E_o = \cup_{i\in\mathcal{I}} E_{o,i}$ the total set of observable events. However, in many situations, the observable events may not be fixed. For instance, communication between agents may lead to an event being observed on occurrence of one transition but not observed on occurrence of a different transition. Also, under energy, bandwidth, or security constraints, a local agent may choose to enable/disable sensors *dynamically* based on its observation history; this also leads to dynamic observations. Thus, in a more general setting, we specify the observations of each agent $i \in \mathcal{I}$ by the mapping $\omega_i : \mathcal{L}(G) \to 2^{E_{o,i}}$. Given an observation mapping, $\omega_i, i \in \mathcal{I}$, we define the projection $P_{\omega_i} : \mathcal{L}(G) \to E_{o,i}^*$ recursively as follows:

$$P_{\omega_i}(\epsilon) = \epsilon, \qquad P_{\omega_i}(s\sigma) = \begin{cases} P_{\omega_i}(s)\sigma & \text{if } \sigma \in \omega_i(s) \\ P_{\omega_i}(s) & \text{if } \sigma \notin \omega_i(s). \end{cases} \tag{1}$$

The inverse of P_{ω_i}, denoted by $P_{\omega_i}^{-1}$, is defined as $P_{\omega_i}^{-1} : E_o{}^* \to 2^{E^{G*}}$ with $P_{\omega_i}^{-1}(s) := \{t \in E^{G*} : P_{\omega_i}(t) = s\}$. The projection P_{ω_i} and its inverse $P_{\omega_i}^{-1}$ are extended to languages in the usual way. Clearly, if the set of observable events is fixed in the sense that $\forall s \in \mathcal{L}(G), \omega_i(s) = E_{o,i}$, then the projection P_{ω_i} reduces to the standard natural projection.

The above definition of observation mapping is language-based; as such, it may require infinite memory to realize. In practice, one is often interested in studying a particular type of dynamic observation, namely, *transition-based* dynamic observation. Formally, for each agent $i \in \mathcal{I}$, we say that an observation mapping $\omega_i : \mathcal{L}(G) \to 2^{E_{o,i}}$ is transition-based if

$$(\forall s, t \in \mathcal{L}(G))[\delta^G(x_0^G, s) = \delta^G(x_0^G, t) \Rightarrow \omega_i(s) = \omega_i(t)]. \tag{2}$$

Thus, a transition-based observation mapping ω_i can also be described by a set of observable transitions $\Omega_i \subseteq TR(G)$ defined by $\Omega_i := \{(x, e) \in TR(G) : \exists s \in \mathcal{L}(G) \text{ s.t. } \delta^G(x_0^G, s) = x \wedge e \in \omega_i(s)\}$.