Proceedings of the 9th International
Workshop on Discrete Event Systems
Göteborg, Sweden, May 28-30, 2008

Decentralized Modular Diagnosis of Concurrent Discrete Event Systems

C. Zhou[♯], Member IEEE, R. Kumar[♯], Fellow IEEE, and R. S. Sreenivas[♭], Member IEEE
[♯]Dept. of Elec. & Comp. Eng., Iowa State Univ., Ames, IA
[♭]Industrial & Enterprise Sys. Eng., Univ. of Illinois, Urbana-Champaign, IL

Abstract— **The problem of decentralized modular fault diagnosis of concurrent discrete event systems, that is composed of a set of component modules, is formulated and studied. In the proposed decentralized modular framework, diagnosis is performed by the local diagnosers, located at the component sites, using their own local observations. This is to ensure the scalability of the approach with respect to the number of component modules, and we require that the local diagnosers be "modularly computable", i.e., their computation should be based on the local models, and not the global models. It is also required that there are no missed-detections (every fault is detected within a bounded number of transitions) and no false-alarms (a fault detection report is issued only when a fault has occurred). We formally define the decentralized modular diagnosis problem and introduce the notion of modular diagnosability as a key property for the existence of desired decentralized modular diagnosers. We show that under this property, the complexity for constructing the local diagnosers is polynomial in the number of local modules. We present a method for testing the modular diagnosability property by reducing it to an instance of a certain codiagnosability property for which known verification techniques exist.**

Keywords: Discrete event systems, concurrent systems, fault diagnosis, failure diagnosis, modularity, modular diagnosability

I. Introduction

All systems are subjected to fault, where a fault is the execution of a behavior that violates the specification of the nominal behaviors. Detection of faults and their isolation is an important exercise. The failure diagnosis problem for discrete event systems (DESs) has been studied in various settings such as centralized untimed setting [13], [18], [17], [8], [29], [28], decentralized/distributed untimed setting [6], [16], [19], [24], [25], [20], [15], [12], [11], [4], [31], in the temporal logic setting [9], in the setting of discrete-time [30] and dense-time [14], [23], [3], [10], and in the probabilistic setting [22], [1].

The work presented in this paper is motivated from failure diagnosis of large scale discrete events systems that are composed of multiple concurrently operating local modules, and are also referred as *concurrent* systems. Examples include communication networks, manufacturing systems, chemical process control, and power systems, all of which are composed of many interacting modules that are physically distributed. For such systems, a centralized failure diagnosis approach is not appropriate (owing to communication and processing delays, and computational scalability), and non-centralized approaches involving multiple local diagnosers have been proposed. These approaches can be classified into decentralized [15], [26], [12], [21] or distributed [6], [19], [16], [2] depending on whether or not the local diagnosers communicate among each other.

For the scalability of a decentralized/distributed diagnosis scheme it is desirable that the complexity of synthesizing local diagnosers grow polynomially in the number of system modules. We refer to this property as *modular computability* and the corresponding decentralized diagnosis problem as *decentralized modular diagnosis* problem. The paper formulates and studies this problem, where in the proposed decentralized modular framework, diagnosis is performed by the local diagnosers, located at the component sites, using their own local observations. We require that the local diagnosers be *modularly computable*, i.e., their computation should be based on the local models, and not the global models. It is also required that there are no missed-detections (every fault is detected within a bounded number of transitions) and no false-alarms (a fault detection report is issued only when a fault has occurred). We introduce the notion of modular diagnosability as a key property for the existence of the desired decentralized modular diagnosers. We show under this property that the complexity for constructing the local diagnosers is polynomial in the number of local modules. We present a method for testing the modular diagnosability property by reducing it to an instance of a certain codiagnosability property. It follows from earlier work on verification of codiagnosability [15] that the verification of modular diagnosability is polynomial in the sizes of the nonfaulty specification and the global plant model.

Modular approaches to fault diagnosis of concurrent DESs have been considered in [7], [5]. In [7], the fault model is assumed to be local (certain events in each local model are faulty, and a global fault of the system is considered to be the execution of a faulty event by a local module). We allow a more general model of a global fault, which is the execution of a trace that *violates a global system specification*. Further, in [7] it is required that a fault be locally detected, i.e., the execution of a faulty event by a module at site-i be detected by the local diagnoser located at the same site-i. We do not impose such a localization requirement for fault-detection, and allow a fault to be detected at any site. Also only a sufficient condition for diagnosability under the said setup is presented in [7]. The set in [5] is similar to that described in [7] with additional restrictions added. For example it is

The research was supported in part by the National Science Foundation under the grants NSF-ECS-0424048, NSF-ECS-0601570 and NSF-ECCS-08013763.

978-1-4244-2593-8/08/$25.00 ©2008 IEEE

required that the diagnosis of a fault at site-i only occur along those traces in which site-i continues to participate (authors refer to this as "persistency of excitation"), whereas in general, there is no apriori guarantee that such a property will hold for a trace executed by a system being diagnosed. Another assumption imposed in [5] is that for all local diagnosers, their *common events are observable*. This is a severe restriction. It is further assumed in [5] that the global plant model is deadlock-free. It is not clear how to relax this assumption without introducing cycles of unobservable events which is disallowed in their framework. Note also that in the context of concurrent systems, even when the local plant modules are deadlock-free, the global plant model can be deadlocking. The decentralized modular diagnosis problem that we formulate and study in this paper has no apriori restrictions.

While the topic of decentralized modular diagnosis of DESs is relatively new, there has been considerable work on decentralized modular control of DESs (see for example [32] and the references there in). Rest of the paper is organized as follows. The notion and preliminaries are given in Section 2. The decentralized modular diagnosis problem and its solution is provided in Section 3. This guarantees the modular computability of the local diagnosers. Section 4 presents a way to verify the existence of decentralized modular diagnosers by establishing connection with codiagnosability. Section 5 concludes the work presented.

II. Notation and Preliminaries

Given an event set Σ, we use $\overline{\Sigma}$ to denote $\Sigma \cup \{\epsilon\}$, and Σ^* to denote the set of all event-traces over Σ, including the trace of zero-length ϵ. A subset $L \subseteq \Sigma^*$ is called a language over Σ. A trace $u \in \Sigma^*$ is a prefix of a trace $v \in \Sigma^*$ if for some trace $w \in \Sigma^*$, $v = uw$. The prefix-closure of $L \subseteq \Sigma^*$, denoted $pr(L)$, is the set of all prefixes of traces in L. L is called prefix-closed or simply closed if $pr(L) = L$. Given a trace $s \in L$, the set of extensions of s in L is the set of traces $\{w \in \Sigma^* \mid sw \in L\}$, denoted $L\backslash s$; s is said to be deadlocking in L if $L\backslash s = \{\epsilon\}$. The length of a trace s, denoted $|s|$, is the number of events in s.

The events executed by a discrete-event system to be diagnosed are observed using sensors, which can be represented using an event observation map, $M : \overline{\Sigma} \to \overline{\Delta}$ satisfying $M(\epsilon) = \epsilon$, where Δ is the set of observed symbols. The observation mask can be extended from events to traces: $M(\epsilon) = \epsilon$ and $\forall s \in \Sigma^*, \sigma \in \Sigma, M(s\sigma) = M(s)M(\sigma)$. The inverse observation mask of M, is a map $M^{-1} : \overline{\Delta} \to 2^{\overline{\Sigma}}$, and can also be extended from observation symbols to set of event traces: $\forall \eta \in \Delta^*$, $M^{-1}(\eta) = \{s \in \Sigma^* \mid M(s) = \eta\}$.

Given the local event sets $\{\Sigma_i\}$ of the modules of a concurrent plant, $\Sigma = \cup_i \Sigma_i$ denotes the set of global events, and $P_i : \Sigma \to \overline{\Sigma}_i$ is used to denote the natural projection from the global event set to the ith local event set. Then the corresponding inverse projection operation is a map: $P_i^{-1} : \overline{\Sigma}_i \to 2^{\overline{\Sigma}}$. Given $\{L_i \subseteq \Sigma_i^*\}$ (representing say the behaviors of the local plant modules), the parallel product of $\{L_i\}$ (representing say the behaviors of the global plant), denoted $||_i L_i$, is defined as:

$$||_i L_i := \{t \in \Sigma^* \mid \forall i : P_i(t) \in L_i\} = \cap_i P_i^{-1}(L_i) \subseteq \Sigma^*.$$

It holds that $s \in \|_i L_i$ if and only if for each i, $P_i(s) \in L_i$. The local observation mask associated with the local observer at site-i is defined as $M_i : \overline{\Sigma}_i \to \overline{\Delta}_i$ satisfying $M_i(\epsilon) = \epsilon$, where Δ_i is the set of locally observed symbols at site-i. Then the inverse of local observation mask is a map: $M_i^{-1} : \overline{\Delta}_i \to 2^{\overline{\Sigma}_i}$. The composition of a local observation mask M_i and a natural projection map P_i gives rise to a new mask function that maps the global events to the locally observed symbols at site-i, $M_i P_i : \overline{\Sigma} \to \overline{\Delta}_i$, satisfying $M_i P_i(\epsilon) = \epsilon$. The inverse of $M_i P_i$ is a map: $(M_i P_i)^{-1} = P_i^{-1} M_i^{-1} : \overline{\Delta}_i \to 2^{\overline{\Sigma}}$.

The notion of separability was introduced in [27] as a key property in a modular framework.

Definition 1: Given the event sets $\{\Sigma_i\}$, a language $K \subseteq \Sigma^*$ is $\{\Sigma_i\}$-separable if $\exists \{K_i \subseteq \Sigma_i^*\}$ such that $K = ||_i K_i$. It is known from [27, Corollary 2.1] that a language $K \subseteq \Sigma^*$ is $\{\Sigma_i\}$-separable if and only if $K = ||_i \ P_i(K)$.

A fault is a deviation from the nominal behaviors. In this paper we use a specification language K to represent the nominal or nonfaulty behaviors. Without loss of generality, K can be assumed to be a sublanguage of the language $L := \|_i L_i \subseteq \Sigma^* = (\cup_i \Sigma_i)^*$ of a concurrent global plant, where for each i, $L_i \subseteq \Sigma_i^*$ is the language of the ith local plant module. Execution of a trace in $L - K$ is the execution of a faulty behavior. Since the prefix of a nonfaulty behavior cannot be a faulty behavior, the nonfaulty specification language K is prefix-closed.

Given a plant and nonfaulty language pair (L, K) defined over an event set Σ, a set of local observation mask functions $\{M_i : \overline{\Sigma}_i \to \overline{\Delta}_i\}$, one for each local diagnoser, the codiagnosability of (L, K) with respect to $\{M_i\}$ is the property under which every faulty trace in $L - K$ can be detected within a bounded delay by some local diagnoser, where the local diagnosers do not communicate with each other.

Definition 2: [15] Consider a plant language $L \subseteq \Sigma^*$, a nonfaulty language $K \subseteq L$, and local observation masks $\{M_i : \overline{\Sigma} \to \overline{\Delta}_i\}$. (L, K) is said to be $\{M_i\}$-codiagnosable if $\exists n \geq 0$ such that

$$\forall s \in L - K, \forall t \in L\backslash s, |t| \geq n \text{ or } st \text{ deadlocks},$$
$$\exists i : M_i^{-1} M_i(st) \cap K = \emptyset.$$

The above property requires the existence of a delay bound n such that for any faulty trace $s \in L - K$ and for any extension $st \in L$ with either $|t| \geq n$ or st deadlocking, exists a local site-i for which it holds that no nonfaulty trace in K is indistinguishable from st (i.e., $M_i^{-1} M_i(st) \cap K = \emptyset$).

Codiagnosability can be verified polynomially in the sizes of L and K [15]. The set of local diagnosers also can be computed polynomially in the size of K [15].

III. Decentralized Modular Diagnosis

In this section we formalize the problem of decentralized modular diagnosis of concurrent DESs, and introduce a

notion of modular diagnosability as a key property under which every fault can be diagnosed using a set of local diagnosers whose computation does not require the construction of the global plant model. Consider the following decentralized modular diagnosis problem: Given a concurrent plant with the set of local behaviors $\{L_i\}$, where $L_i \subseteq \Sigma_i^*$ is the language generated by the ith plant module with event set $\{\Sigma_i\}$, a prefix-closed specification language $K \subseteq L = \|_i L_i$ representing the set of nonfaulty behaviors, and a set of local observation masks $\{M_i : \overline{\Sigma}_i \rightarrow \overline{\Delta}_i\}$, determine whether there exists a set of local diagnosers such that the following requirements are met:

- There are no missed-detections, i.e., any faulty trace $s \in ||_i L_i - K$ can be detected by some local diagnoser in a bounded delay, say $n \geq 0$, (i.e., within the occurrence of a trace $t \in (||_i L) \backslash s$ such that either $|t| \geq n$ or st is deadlocking);
- There are no false-alarms, i.e., any nonfaulty trace is never reported as faulty; and
- The local diagnosers are *modularly computable*, i.e., the computation of the ith local diagnoser depends only on the local plant model L_i, the local observation mask M_i, and the nonfaulty specification K.

The ith local diagnoser observes the locally executed traces filtered through its local observation mask, i.e., traces in $M_i(L_i)$, and issues a fault detection decision, either "sure" (denoted 1), or "unsure", denoted ϕ. In other words, the ith local diagnoser is a map, $D_i : M_i(L_i) \rightarrow \{1, \phi\}$ (equivalently, it is a map, $D_i : M_i(P_i(L)) \rightarrow \{1, \phi\}$). The three requirements mentioned above can be formalized as follows:

- There are no missed-detections, i.e., $\exists n \geq 0$:

$$\forall s \in L - K, \forall t \in L \backslash s : |t| \geq n \text{ or } st \text{ deadlocks}, \exists i : D_i(M_i P_i(st)) = 1.$$

- There are no false-alarms, i.e.,

$$\forall s \in K, \forall i : D_i(M_i P_i(s)) \neq 1.$$

- Local diagnosers are modularly computable, i.e., for each i, the computation of D_i only depends on L_i, M_i, P_i, K (and not on $L = \|_i L_i$).

Next we present a necessary and sufficient condition for the solution of the decentralized modular diagnosis problem formulated above. It turns out that the solution requires the existence of "local nonfaulty specifications" $\{K_i \subseteq L_i\}$ with the following properties:

(C1) Each global fault can be locally detected within a finite bounded delay (so that there are no missed-detections): $\exists n \geq 0$:

$$\forall s \in L - K, \forall t \in L/s : |t| \geq n \text{ or } st \text{ deadlocks}, \exists i : M_i^{-1} M_i(P_i(st)) \cap K_i = \emptyset.$$

(C2) Each global nonfaulty trace is locally indistinguishable from local nonfaulty traces (so that there are no false-alarms):

$$\forall s \in K, \forall i : M_i^{-1} M_i(P_i(s)) \cap K_i \neq \emptyset.$$

Accordingly we define the following notion of modular diagnosability.

Definition 3: Consider a plant language $L = ||_i L_i$, where $L_i \subseteq \Sigma_i^*$, a nonfaulty language $K \subseteq L$, local observation masks $\{M_i : \overline{\Sigma}_i \rightarrow \overline{\Delta}_i\}$, and projection maps $\{P_i : \overline{\Sigma} \rightarrow \overline{\Sigma}_i\}$. (L, K) is $\{M_i, \Sigma_i\}$-*modularly diagnosable* if there exist modularly computable local languages $\{K_i \subseteq L_i\}$ such that C1 and C2 hold.

Note Definition 3 allows the global and local plant models to be deadlocking. Thus unlike [7], [5] we do not apriori assume that the global plant model is deadlock-free. This is because even when the local plant models are deadlock-free, the global plant model can be deadlocking. The following example illustrates Definition 3.

Example 1: Consider a scenario in which two computer users access two data buffers shown in Figure 1. The two computer users are physically separated and the two buffers are located in a third computer. Suppose a data can only be accessed by one user at a time and otherwise, an error occurs. Each user can access each data buffer. The events that user-a (resp., user-b) accesses buffers 1 and 2 are denoted as a_1 (resp., b_1) and a_2 (resp., b_2), respectively. Assume the local diagnoser on each computer does not track the buffer addresses, and thus, a_1 and a_2 (resp., b_1 and b_2) are locally indistinguishable. Suppose the simultaneous access of the buffer 1 (resp., 2), denoted c_1 (resp., c_2), can be reported as error event to diagnoser-a (resp., diagnoser-b). Then c_1 (resp., c_2) is observable for diagnoser-a (resp., diagnoser-b) and unobservable for diagnoser-b (resp., diagnoser-a).

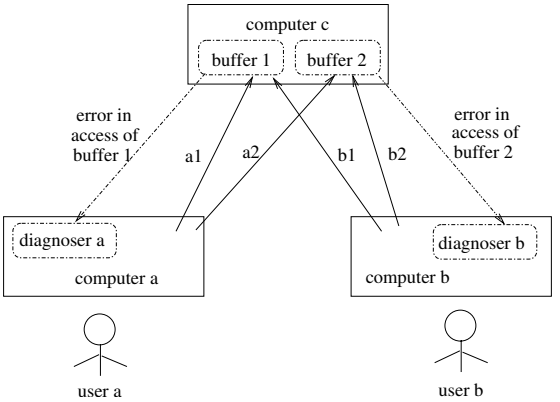

Fig. 1. Two users access two buffers

The scenario can be modeled by a plant language in which any simultaneous access of any of the buffers is followed by an error event,
$L = pr(a_1 b_1 c_1 + b_1 a_1 c_1 + a_2 b_2 c_2 + b_2 a_2 c_2 + a_2 b_1 + b_1 a_2 + a_1 b_2 + b_2 a_1)$, and their local projections (which correspond to the local plant models) are given by,
$L_1 = pr(a_1 c_1 + a_2 c_2) \subseteq \Sigma_1^*$, where $\Sigma_1 = \{a_1, a_2, c_1, c_2\}$, and
$L_2 = pr(b_1 c_1 + b_2 c_2) \subseteq \Sigma_2^*$, where $\Sigma_2 = \{b_1, b_2, c_1, c_2\}$.
Then traces in L not extended by an error event are nonfaulty and so the nonfaulty language is given by,

$$K = pr(a_1 b_2 + b_2 a_1 + a_2 b_1 + b_1 a_2).$$

Then $L - K = \{a_2b_2, a_2b_2c_2, b_2a_2, b_2a_2c_2, a_1b_1, a_1b_1c_1, b_1a_1, b_1a_1c_1\}$ consists of the set of faulty traces. The local observation masks are given by:

$$M_1(a_1) = M_1(a_1) = a, M_1(c_1) = c_1, M_1(c_2) = \epsilon,$$
$$M_2(b_1) = M_2(b_2) = b, M_2(c_1) = \epsilon, M_2(c_2) = c_2.$$

We show that the modular diagnosability holds by choosing $K_1 := pr(a_1)$ and $K_2 := pr(b_2)$. To check C1 of Definition 3, set $n = 1$. Then for

- $s = a_2b_2 \in L - K, t = c_2 \in L\backslash s$ with $|t| \geq 1, M_2^{-1}M_2(P_2(st)) \cap K_2 = \{c_1^*(b_1 + b_2)c_1^*c_2c_1^*\} \cap K_2 = \emptyset$;
- $s = a_1b_1 \in L - K, t = c_1 \in L\backslash s$ with $|t| \geq 1, M_1^{-1}M_1(P_1(st)) \cap K_1 = \{c_2^*(a_1 + a_2)c_2^*c_1c_2^*\} \cap K_1 = \emptyset$.

One can verify that (C1) holds.

Next it can also be verified that C2 holds since for $s = \epsilon \in K$, $M_i^{-1}M_i(P_i(s)) \cap K_i = \{\epsilon\}$ $(i = 1, 2)$; $\forall s \in K - \{\epsilon\}$, $M_1^{-1}M_1(P_1(s)) \cap K_1 = \{a_1\}$, and $M_2^{-1}M_2(P_2(s)) \cap K_2 = \{b_2\}$. Thus, (L, K) is $\{M_i, \Sigma_i\}$-modularly diagnosable.

The following theorem provides a condition under which the decentralized modular diagnosis problem is solvable.

Theorem 1: Consider a plant language $L = ||_i L_i$, where $L_i \subseteq \Sigma_i^*$, a nonfaulty language $K \subseteq L$, local observation masks $\{M_i : \overline{\Sigma}_i \to \overline{\Delta}_i\}$, and projection maps $\{P_i : \Sigma \to \overline{\Sigma}_i\}$. The following are equivalent:

I. There exist modularly computable diagnosers $\{D_i : M_i(L_i) \to \{1, \phi\}\}$ satisfying
(No Missed-Detection)
$\exists n \geq 0$: $\forall s \in L - K$, $\forall t \in L\backslash s$ with $|t| \geq n$ or st deadlocks, $\exists i$: $D_i(M_iP_i(st)) = 1$,
and
(No False-Alarm)
$\forall s \in K, \forall i : D_i(M_iP_i(s)) \neq 1$.

II. (L, K) is $\{M_i, \Sigma_i\}$-modularly diagnosable.

Proof: (I $\Leftarrow$ II) Given modularly computable $\{K_i \subseteq L_i\}$ such that C1 and C2 hold, for each i and for each $u_i \in M_i(L_i)$ define:

$$D_i(u_i) := \begin{cases} 1 & \text{if } M_i^{-1}(u_i) \cap K_i = \emptyset \\ \phi & \text{otherwise.} \end{cases}$$

Then since K_i is modularly computable, so is D_i.

We first show that C1 implies (No Missed-Detection), by showing $\exists n \geq 0$: $\forall s \in L - K$, $\forall t \in L\backslash s$ with $|t| \geq n$ or st deadlocks, $\exists i : D_i(M_iP_i(st)) = 1$. We claim that n can be chosen to the same as that appearing in C1. From C1 for $s \in L - K, st \in L\backslash s$ with $|t| \geq n$ or st deadlocks, $\exists i : M_i^{-1}M_i(P_i(st)) \cap K_i = \emptyset$. Letting $u_i = M_iP_i(st)$, it follows from the definition of D_i that $1 = D_i(u_i) = D_i(M_iP_i(st))$ as desired.

Next we show that C2 implies (No False-Alarm), i.e., $\forall s \in K, \forall i$: $D_i(M_iP_i(s)) \neq 1$. For all i, define $u_i := M_iP_i(s)$. Then from C2, $M_i^{-1}(u_i) \cap K_i \neq \emptyset$. It follows from the definition of D_i that for all i, $1 \neq D_i(u_i) = D_i(M_iP_i(s))$ as desired.

(I $\Rightarrow$ II) Given local diagnosers $\{D_i\}$, we define for each i,

$$K_i := \{u_i \in L_i \mid D_i(M_i(u_i)) \neq 1\}.$$

Then since D_i is modularly computable, so is K_i.

We first show that (No Missed-Detection) implies C1, i.e., $\exists n \geq 0 : \forall s \in L - K, \forall t \in L\backslash s$ with $|t| \geq n$ or st deadlocks, $\exists i : M_i^{-1}M_iP_i(st) \cap K_i = \emptyset$. We claim that n can be chosen to be the same as that appearing in (No Missed-Detection). From (No Missed-Detection) for $s \in L - K, t \in L\backslash s$ with $|t| \geq n$ or st deadlocks, $\exists i : D_i(M_iP_i(st)) = 1$. We show that $M_i^{-1}M_i(P_i(st)) \cap K_i = \emptyset$. Suppose for contradiction that exists $u_i \in K_i$ such that $M_i(u_i) = M_iP_i(st)$. Then $D_i(M_iP_i(st)) = D_i(M_i(u_i)) = 1$. Then from the definition of K_i, $u_i \notin K_i$, which is a contradiction.

Next we show that (No False-Alarm) implies C2, i.e., $\forall s \in K, \forall i : M_i^{-1}M_iP_i(s) \cap K_i \neq \emptyset$. Pick $s \in K$. Then from (No False-Alarm), $\forall i : D_i(M_iP_i(s)) \neq 1$. Then for all i, $P_i(s) \in K_i$. It follows that for all i, $M_i^{-1}M_iP_i(s) \cap K_i \neq \emptyset$. This completes the proof. ■

We now make a few observations about the modular diagnosability property.

Remark 1: It is important to note that separability of the global nonfaulty specification is not necessary for modular diagnosability. Consider the setting of Example 1. Then as concluded in the example, (L, K) is modularly diagnosable. However K is not separable: For example $a_2b_2, b_2a_2 \in ||_iP_i(K) - K$. This is in contrast to the setting of decentralized modular control where the separability serves as a necessary condition [32].

In Example 1 we chose, $\{K_i\}$ such that $||_iK_i = pr(a_1b_2 + b_2a_1) \subseteq K$. This, however, is not required. In fact one can verify that C1 and C2 also hold for $\{K_1 := pr(a_1 + a_2), K_2 := pr(b_1 + b_2)\}$, and in which case $||_iK_i \not\subseteq K$. Note this example also illustrates that the choice of $\{K_i\}$ required to establish modular diagnosability need not be unique.

Finally, as shown via Remark 3 below, modular diagnosability is not comparable to "local diagnosability".

IV. Test for Modular Diagnosability

The definition of modular diagnosability is existential that requires the existence of modularly computable languages $\{K_i \subseteq L_i\}$ satisfying C1 and C2. C1 can be trivially satisfied by simply choosing $K_i = \emptyset$. This however will violate C2. On the other hand C2 can be trivially satisfied by simply choosing $K_i = L_i$, which however will violate C1. Thus K_i can neither be chosen to be too small (there will be false-alarms) or too large (there will be missed-detections). The selection of an appropriate K_i is discussed in this section, which also serves to provide a way to verify modular diagnosability.

We show that the modular diagnosability of (L, K) is equivalent to the *codiagnosability* [15] of (L, K) with respect to the local mask functions $\{M_iP_i\}$. Note in the modular setting, the ith local diagnoser observes local events in Σ_i filtered through the local mask M_i. This is equivalent to observing global events in Σ filtered through the composed

mask M_iP_i. Thus the codiagnosability of (L, K) with respect to the masks $\{M_iP_i\}$ is an expected necessary condition. That it is also a sufficient condition is an interesting observation since the codiagnosability only requires that each fault be detected by some local diagnoser within a bounded delay (and it does not require that each local diagnoser be modularly computable). We are able to show that the modular computability is guaranteed whenever the codiagnosability property holds. In fact we show that K_i can be chosen to be the same as $P_i(K)$, which is the key observation behind the next theorem.

It is important however to point out the differences between the modular and the decentralized diagnosis frameworks: (i) Only a global plant model is required in the decentralized setting, where the modular setting requires local plant models, (ii) The mask functions are global in the decentralized setting (and hence the diagnosers are defined over the global plant behaviors) whereas they are local in the modular setting (and hence the diagnosers are defined over the local plant behaviors).

Theorem 2: Consider a plant language $L = ||_i L_i$, where $L_i \subseteq \Sigma_i^*$, a nonfaulty language $K \subseteq L$, local observation masks $\{M_i : \overline{\Sigma}_i \rightarrow \overline{\Delta}_i\}$, and projection maps $\{P_i : \Sigma \rightarrow \overline{\Sigma}_i\}$. (L, K) is $\{M_i, \Sigma_i\}$-modularly diagnosable if and only if (L, K) is $\{M_iP_i\}$-codiagnosable.

Proof: ($\Leftarrow$) For each i, define $K_i := P_i(K)$. Then $K_i = P_i(K) \subseteq P_i(L) \subseteq L_i$. Also clearly K_i is modular computable (its computation does not require the computation of the global plant model L). We first show C2 holds, i.e., $\forall s \in K, \forall i$: $M_i^{-1}M_i(P_i(s)) \cap K_i \neq \emptyset$. This follows from the fact that $P_i(s) \in P_i(K) = K_i$ and $P_i(s) \in M_i^{-1}M_i(P_i(s))$, and so $P_i(s) \in M_i^{-1}M_i(P_i(s)) \cap K_i \neq \emptyset$.

Next we show C1 holds, i.e., $\exists n \geq 0 : \forall s \in L - K, \forall t \in L\backslash s : |t| \geq n$ or st deadlocks, $\exists i : M_i^{-1}M_i(P_i(st)) \cap K_i = \emptyset$. We claim that n can be chosen to be the same as that appearing in the definition of codiagnosability. Pick $s \in L - K, t \in L\backslash s : |t| \geq n$ or st deadlocks. Then from codiagnosability, exists i such that $(M_iP_i)^{-1}M_iP_i(st) \cap K = \emptyset$. We claim that this implies $M_i^{-1}M_i(P_i(st)) \cap K_i = \emptyset$. Otherwise exists $u_i \in K_i = P_i(K)$ such that $M_i(u_i) = M_iP_i(st)$. Since $u_i \in P_i(K)$, there exists $u \in K$ such that $P_i(u) = u_i$. Then $M_iP_i(u) = M_i(u_i) = M_iP_i(st)$. Thus $u \in K \cap (M_iP_i)^{-1}M_iP_i(st)$, a contradiction.

($\Rightarrow$) We need to show that $\exists n \geq 0$: $\forall s \in L - K$, $t \in L\backslash s$, $|t| \geq n$ or st deadlocks, $\exists i$: $(M_iP_i)^{-1}M_iP_i(st) \cap K = \emptyset$. We claim that n can be chosen to be the same as that appearing in C1 of modular diagnosability. Pick $s \in L - K, t \in L\backslash s : |t| \geq n$ or st deadlocks. Then from C1 exists i such that $M_i^{-1}M_i(P_i(st)) \cap K_i = \emptyset$. We claim that this implies $(M_iP_i)^{-1}M_iP_i(st) \cap K = \emptyset$. Otherwise exists $u \in K$ such that $M_iP_i(u) = M_iP_i(st)$. Since $u \in K$, from C2 there exists $u_i \in K_i$ such that $M_i(u_i) = M_iP_i(u)$. Then $M_i(u_i) = M_iP_i(u) = M_iP_i(st)$. Thus $u_i \in K_i \cap M_i^{-1}M_i(P_i(st))$, a contradiction. ■

Remark 2: Theorem 2 is able to shed interesting insights into the nature of the decentralized modular diagnosis problem. For example it establishes that whenever the languages $\{K_i \subseteq L_i\}$ satisfying the conditions of modular diagnosability exist, the languages $\{P_i(K)\}$ also work. Further, the $\{M_i, \Sigma_i\}$-modular diagnosability is equivalent to $\{M_iP_i\}$-codiagnosability, i.e., whenever the diagnosis can be performed in the decentralized setting under the masks $\{M_iP_i\}$, it can also be performed in the modular setting.

It follows from the earlier work on synthesis of diagnosers in the decentralized setting [15] that the local diagnosers can be chosen to be the generators of the languages $\{M_iP_i(K)\}$, which are computable linearly in the size of the nonfaulty language K. It follows that the local diagnosers are modularly computable. It further follows from the earlier work on verification of codiagnosability [15] that modular diagnosability can be verified polynomially in the sizes of L_i (for each i) and K, and exponentially in the number of local modules. (Note the verification, being an off-line computation, is not required to be performed on-line in real-time.) The modular computability of the local diagnosers, on the other hand, guarantees their on-line real-time implementability.

We showed that under modular diagnosability $\{P_i(K))\}$ can be used as the local nonfaulty specifications. Then we can have the notion of local diagnosability of $(L_i, P_i(K))$ with respect to the local mask M_i. It would be possible to check modular diagnosability without having to compute the global plant model if it so happens that the local diagnosability implies modular diagnosability. This however is not true as illustrated in the following remark which establishes a stronger assertion that the notion of local and modular diagnosabilities are not comparable.

Remark 3: We show that

$$\forall i : (L_i, P_i(K)) \ M_i\text{-diagnosable}$$
$$\not\Leftrightarrow \quad (L, K) \ \{M_i, \Sigma_i\}\text{-modularly diagnosable.}$$

Consider $L = L_1||L_2$ with $L_1 = pr(a_1c_1^* + a_2c_2^*)$ and $L_2 = pr(b_1c_1^* + b_2c_2^*)$, where $\Sigma_1 = \{a_1, a_2, c_1, c_2\}$ and $\Sigma_2 = \{b_1, b_2, c_1, c_2\}$. Assume $M_1(a_1) = M_1(a_2) = a$, $M_2(b_1) = M_2(b_2) = b$, $M_i(c_1) = c_1$ and $M_i(c_2) = c_2$ for $i = 1, 2$. Let $K = pr(a_2b_1 + b_1a_2)$. Then for $i = 1, 2$, $(L_i, P_i(K))$ is M_i-diagnosable. However, the faulty trace $a_1b_2 \in L - K$ cannot be detected at any local site since $a_2b_1 \in (M_1P_1)^{-1}M_1P_1(a_1b_2) \cap K \neq \emptyset$ and $b_1a_2 \in (M_2P_2)^{-1}M_2P_2(a_1b_2) \cap K \neq \emptyset$. Thus, (L, K) is not $\{M_iP_i\}$-codiagnosable. From Theorem 2, (L, K) is not $\{M_i, \Sigma_i\}$-modularly diagnosable.

For the converse, consider Example 1 which we showed to be modularly diagnosable. It can be verified that $(L_i, P_i(K))$ is not M_i-diagnosable for $i = 1, 2$: For the faulty trace $a_2c_2 \in L_1 - P_1(K)$ it holds that $a_2 \in M_1^{-1}M_1(a_2c_2) \cap P_1(K) \neq \emptyset$, and for the faulty trace $b_1c_1 \in L_2 - P_2(K)$ it holds that $b_1 \in M_2^{-1}M_2(b_1c_1) \cap P_2(K) \neq \emptyset$.

V. Conclusion

We formulated and studied the problem of decentralized modular diagnosis of concurrent (large-scale) DESs. Besides correctness (no missed-detections and no false-alarms), the modular computability of the local diagnosers is required for

the approach to be scalable. (Modular computability requires the computation must avoid computing the global plant model which can be prohibitive for large-scale systems.) We showed that the solution requires the existence of modularly computable local nonfault specifications satisfying certain properties. We refereed to this property as *modular diagnosability*, and provided a method for verifying it by reducing to an instance of a codiagnosability property. We also showed that whenever the decentralized modular diagnosis problem admits a solution, the local projections of the global nonfault specification can be used as the local nonfault specifications. We illustrated that the choice for the solution is not unique in general. We also noted that the existence of a solution does not require the separability of the nonfault specification, whereas in contrast, this is a necessary condition for the existence of a decentralized modular control. We also showed that the modular diagnosability property is incomparable with the local diagnosability property. Finding a sufficient condition under which the latter implies the former is a direction for future research as this will allow checking modular diagnosability without having to compute the global plant model. (Note that the computation of the local diagnosers avoids computing the global plant model, but checking their existence does. We believe the reason is that the existence problem is computationally hard, and establishing this fact formally is a future research task.)

REFERENCES

[1] E. Athanasopoulou and C. N. Hadjicostis. Probabilistic approaches to fault detection in networked discrete event systems. *IEEE Transactions on Neural Networks*, 16(5):1042–1052, 2005.

[2] R. K. Boel and J. H. van Schuppen. Decentralized failure diagnosis for discrete-event systems with constrained communication between diagnosers. In *Proceedings of International Workshop on Discrete Event Systems*, 2002.

[3] P. Bouyer, F. Chevalier, and D. D'Souza. Fault diagnosis using timed automata. In *Proceeding of the 8th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'05)*, Edinburgh, 2005.

[4] O. Contant, S. Lafortune, and D. Teneketzis. Diagnosis of intermittent faults. *Discrete Event Dynamical Systems: Theory and Application*, 14:171–202, 2004.

[5] O. Contant, S. Lafortune, and D. Teneketzis. Diagnosability of discrete event systems with modular structure. *Discrete Event Dynamic Systems: Theory and Applications*, 16:9–37, 2006.

[6] R. Debouk, S. Lafortune, and D. Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete event systems. *Discrete Event Dynamical Systems: Theory and Applications*, 10:33–79, 2000.

[7] R. Debouk, R. Malik, and B. Brandin. A modular architecture for diagnosis of discrete event systems. In *Proc. of IEEE Conf. on Decision and Control*, pages 417–422, Las Vegas, NV, December 2002.

[8] S. Jiang, Z. Huang, V. Chandra, and R. Kumar. A polynomial time algorithm for diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 46(8):1318–1321, 2001.

[9] S. Jiang and R. Kumar. Failure diagnosis of discrete event systems with linear-time temporal logic fault specifications. *IEEE Transactions on Automatic Control*, 49(6):934–945, 2004.

[10] S. Jiang and R. Kumar. Diagnosis of dense-time systems using digital-clocks. In *Proceedings of the 25th American Control Conference*, pages 6051–6056, Minneapolis, MN, June 2006.

[11] S. Jiang, R. Kumar, and H. E. Garcia. Diagnosis of repeated/intermittent failures in discrete event systems. *IEEE Transactions on Robotics and Automation*, 19(2):310–323, 2003.

[12] R. Kumar and S. Takai. Inference-based ambiguity management in decentralized decision making: Decentralized failure diagnosis of discrete event systems. In *Proceedings of the 25th American Control Conference*, pages 6069–6074, Minneapolis, MN, June 2006.

[13] F. Lin. Diagnosability of discrete event systems and its applications. *Discrete Event Dynamic Systems: Theory and Applications*, 4(1):197–212, 1994.

[14] D. Pandalai and L. Holloway. Template languages for fault monitoring of timed discrete event processes. *IEEE Transactions on Automatic Control*, 45(5):868–882, May 2000.

[15] W. Qiu and R. Kumar. Decentralized failure diagnosis of discrete event systems. *IEEE Transactions on Systems, Man, and Cybernetics—A*, 36(2):384–395, 2006.

[16] S. L. Ricker and J. H. van Schuppen. Decentralized failure diagnosis with asynchronous communication between supervisors. In *Proceedings of the European Control Conference*, pages 1002–1006, 2001.

[17] M. Sampath and S. Lafortune. Active diagnosis of discrete event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, 1998.

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

[19] R. Sengupta and S. Tripakis. Decentralized diagnosis of regular language is undecidable. In *Proceedings of IEEE Conference on Decision and Control*, pages 423–428, Las Vegas, NV, December 2002.

[20] R. Su and W. M. Wonham. Global and local consistencies in distributed fault diagnosis for discrete-event systems. *IEEE Transactions on Automatic Control*, 50(12):1923–1935, 2005.

[21] S. Takai and R. Kumar. Decentralized diagnosis for nonfailures of discrete event systems using inference-based ambiguity management. In *Proceeding of 2006 International Workshop on Discrete Event Systems*, pages 242–247, Ann Arbor, MI, July 2006.

[22] D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Transactions on Automatic Control*, 50(4):476–498, 2005.

[23] S. Tripakis. Fault diagnosis for timed automata. In *Formal Techniques in Real Time and Fault Tolerant Systems*, volume 2469 of *Lecture Notes in Computer Science*. Springer Verlag, 2002.

[24] Y. Wang, T.-S. Yoo, and S. Lafortune. New results on decentralized diagnosis of discrete-event systems. In *Proceedings of 2004 Annual Allerton Conference*, 2004.

[25] Y. Wang, T.-S. Yoo, and S. Lafortune. Decentralized diagnosis of discrete event systems using unconditional and conditional decisions. In *Proceedings of the 44th IEEE Conference on Decision and Control and 2005 European Control Conference*, pages 6298–6304, Seville, Spain, December 2005.

[26] Y. Wang, T.-S. Yoo, and S. Lafortune. Diagnosis of discrete event systems using decentralized architectures. *Discrete Event Dynamic Systems: Theory and Applications*, 17(2):233–263, 2007.

[27] Y. Willner and M. Heymann. Supervisory control of concurrent discrete-event systems. *International Journal of Control*, 54(5):1143–1169, 1991.

[28] T. S. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *IEEE Transactions on Automatic Control*, 47(9):1491–1495, 2002.

[29] S. H. Zad, R. H. Kwong, and W. M. Wonham. Fault diagnosis in discrete-event systems: Framework and model reduction. *IEEE Transactions on Automatic Control*, 48(7):1199–1212, 2003.

[30] S. H. Zad, R. H. Kwong, and W. M. Wonham. Fault diagnosis in discrete-event systems: Incorporating timing information. *IEEE Transactions on Automatic Control*, 50(7):1010–1015, 2005.

[31] C. Zhou and R. Kumar. Computation of diagnosable fault-occurrence indices for systems with repeatable-faults. In *Proceedings of the 44th IEEE Conference on Decision and Control and 2005 European Control Conference*, pages 6311–6316, Seville, Spain, December 2005.

[32] C. Zhou, R. Kumar, and R. S. Sreenivas. Decentralized modular control of concurrent discrete event systems. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 5918–5923, New Orleanes, LA, December 2007.