DIAGNOSIS AND COMMUNICATION IN DISTRIBUTED SYSTEMS

Raja Sengupta*
California PATH, University of California at Berkeley
1357 S. 46th Street,# 452, Richmond, CA 94804
raja@path.berkeley.edu

Keywords: diagnosis, discrete event systems, distributed systems, protocols

July 26, 1998

Abstract

This paper discusses diagnosis problems in distributed systems within the context of a language-theoretic discrete event formalism. A distributed system is seen as a system with multiple spatially separated sites with each site having a diagnoser that observes some of the events generated by the system and diagnoses the faults associated with the site. We allow the diagnosers to share information by sending messages to each other. The existence and synthesis of diagnosers is investigated. The formulation and results are motivated by the diagnosis of failures in a wireless LAN.

1 Introduction

We are interested in understanding the design of diagnostics for distributed systems. This theoretical work is motivated by our experience with the design of distributed diagnostics for coordinating vehicle systems [5, 10] and wireless local area networks [3, 6]. These systems are comprised of spatially separated sites (e.g., vehicles or radios) of semi-autonomous activity. Since these systems operate under distributed control, it is desirable that each site be able to diagnose (detect and isolate) its own failures. In general, we find that a site requires information from other sites to isolate some of its failures. This sharing of information is realized through protocols executed over a communication network. We seek insight into the role of communication in diagnostic design for distributed systems.

The relevant literature is as follows. Fault detection in distributed discrete event systems has been investigated in the context of Petri Net models with applications to telecommunication networks in [2] and [1]. Similar problems have been studied in the context of template languages for fault monitoring in [9]. Within the context of automata and language-theoretic models, centralized diagnostic design for discrete event systems was investigated in [12, 13]. The results presented in this paper are also in the context of a language-theoretic discrete event formalism and generalize some of the results in [12].

The system for which diagnostics is to be designed is called the *plant.* Diagnostic design for a distributed plant entails the design of several communicating diagnostic processes (hereafter *diagnosers*). Each diagnoser observes some of the events generated by the plant. Failure events are assumed to be unobservable to all the diagnosers. Each diagnoser tracks the plant by observing some subsequence of the event sequence executed by the plant, communicates with other diagnosers through the generation and reception of message events, infers the normal or failed status of the system usings its observations and communications, isolates the particular failures, if any, that have occurred, and generates failure messages reporting the isolated failures to the controllers. Figure 1 shows this architecture. The diagnosers (e.g., D_1, D_2, D_3 in figure 1) are assumed to communicate over a communication system distinct from the plant. A distributed plant for which there exist diagnosers that are able to diagnose the failures at the respective sites, is said to be a *diagnosable* system. This paper formalizes the distinction between diagnosable and non-diagnosable systems.

We shall concern ourselves with three kinds of diagnosability, i.e., a distributed system may be *independently diagnosable, decentrally diagnosable,* and *centrally diagnosable.* We think of a distributed system as being independently diagnosable if for every site, the failures of interest to the site can be diagnosed using local information alone, i.e., without information from other sites. A distributed system is decentrally diagnosable if there exists an inter-diagnoser communication scheme providing remote information that together with local observation will provide each site with enough information to diagnose the failures of interest. Finally,

*Research supported in part by Office of Naval Research grant 442427-25828 and CALTRANS PATH MOU 331.

we think of a distributed system as being centrally diagnosable if all the failures at all sites can be diagnosed by a central observer observing instantaneously any event observed at any site. Our intution suggests that idependently diagnosable systems should be decentrally diagnosable and decentrally diagnosable systems should be centrally diagnosable. However, the converse may not be true. We seek to understand when centrally diagnosable systems fail to be decentrally diagnosable, i.e., why no inter-diagnoser messaging scheme exists even when the global observation structure is rich enough to support diagnosis.

The plant is formalized as a discrete event system that in general *concurrently, spontaneously* and *asynchronously* executes multiple processes. However, the representation of concurrency is restricted to an interleaved semantics. The behavior of a single process is an event sequence. The behavior of a collection of processes is also an event sequence. Accordingly, the set of possible behaviors of the plant is a language $L_p \subseteq \Sigma_p^*$, where Σ_p is a finite alphabet. We assume the diagnosers have no control capabilities with respect to the plant. For example, in the language of supervisory control ([11]), all plant events appear uncontrollable to the diagnosers. The diagnosers may not force, disable or delay the execution of plant events. Diagnosers communicate by synchronizing on message events. We think of these messages as being sent over a faultless communication system distinct from the plant. Therefore the generation of message and plant events is completely asynchronous. No assumptions are made to bound delays in the inter-diagnoser communication system. Therefore in the interleaved semantics, arbitrarily many plant events may preceed or succeed a diagnoser message event.

The structure of the paper is as follows. Section 2 formulates the problem mathematically. Section 3 presents a motivating example. Section 4 presents results characterising diagnosable and non-diagnosable systems. Section 5 summarizes the results of the paper.

2 Problem Formulation

In this section we define the entities in figure 1, and the three kinds of diagnosability.

As stated in the introduction, the plant is modeled as a language L_p that lies in the kleene closure ${\Sigma_p}^*$, of a finite alphabet Σ_p. $\mathcal{A} = \{1, \ldots, |\mathcal{A}|\}$ is an index set of sites, $|\mathcal{A}|$ is the cardinality of the set $\mathcal{A}$ and the number of sites, $\Sigma_{poi} \subseteq \Sigma_p$, is the set of plant events observed at site i, and $\Sigma_{fi} \subseteq \Sigma_p$, is the set of plant failure events that should be diagnosed by site i. Let $\cup_{i=1}^{i=|\mathcal{A}|}\Sigma_{poi} = \Sigma_{po}$ and $\cup_{i=1}^{i=|\mathcal{A}|}\Sigma_{fi} = \Sigma_f$ be respectively the set of all observable and failure events in the plant It is assumed that

$$\Sigma_f \cup \Sigma_{po} = \emptyset,$$

i.e., all failures are unobservable to all sites. We assume that L_p is prefix-closed and live. $\Sigma_{puo} = \Sigma_p - \Sigma_{po}$, denotes the set of unobservable plant events.

The task of each diagnoser is to process observations and generate failure messages that isolate the failures that have occurred for the benefit of fault management control. Let Σ_{mfi} denote the set of failure messages generated by site i and $\Sigma_{mf} = \cup_{i=1}^{i=|\mathcal{A}|}\Sigma_{mfi}$ be the set of all failure messages. Without loss of generality we pick Σ_{mf} to be any set of symbols, disjoint from all the other alphabets, that can be put in one-to-one correpondence with Σ_f. Let $\theta : \Sigma_f \rightarrow \Sigma_{mf}$ be a bijection. Then $\theta[\Sigma_{fi}] = \Sigma_{mfi}$.

We let Σ_{mi} denote the set of inter-diagnoser communication messages generated by site i. It is assumed that the Σ_{mi} are pairwise disjoing and $\Sigma_m = \uplus_{i=1}^{i=|\mathcal{A}|}\Sigma_{mi}$, is disjoint from Σ_p amd Σ_{mf}. The alphabet $\Sigma = \Sigma_p \uplus \Sigma_m \uplus \Sigma_{mf}$, is the set of all events generated by the plant and the diagnosers.

The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ specifies the plant, the decentralized observation structure, and the failures to be diagnosed by each diagnoser. Since the concurrent operation of the plant and diagnosers is represented in an interleaved semantics, their combined behavior is an event trace in Σ^*. Let $L \subseteq \Sigma^*$ denote the set of all possible interleaved behaviors of the plant and diagnosers. Thus L specifies a design, i.e., the generation of inter-diagnoser communication messages, and the generation of failure messages by the diagnosers. We will refer to L as the designed language. Our objective is to design L and the associated $\{\Sigma_{mi}\}_{i\in\mathcal{A}}$ so that the diagnosers will generate failure messages when failures have occurred and not generate such messages when failures have not occurred.

2.1 Admissible Designs

We require that diagnosers not force, disable, or delay the generation of plant events. We will also require that designs be causal, i.e., the messages generated by a diagnoser should be a function of its observations and communications. Designs having these properties are said to be admissible and formalized as follows. Let P_A denotes the projection of a trace in Σ^* onto A^*, for an alphabet $A \subseteq \Sigma$.

Definition 1 *L is an admissible design iff*

1. *L is prefix-closed.*
2. *L is plant consistent, i.e., $L_p \subseteq L, P_{\Sigma_p}(L) = L_p, (\forall w \in L, P_{\Sigma_p}(w)\sigma \in L_p \Rightarrow w\sigma \in L)$.*
3. *L is causal, i.e., $\forall i \in \mathcal{A}, \sigma \in \Sigma_{mfi} \cup \Sigma_{mi}, u, v \in L,$*

$$(u\sigma \in L) \wedge (P_{\Sigma_{o_i}}(u) = P_{\Sigma_{o_i}}(v)) \Rightarrow v\sigma \in L.$$

The plant consistency assumption says that the diagnosers cannot force, disable, or delay the behavior of the plant. The causality condition says that

the messaging scheme designed for a diagnoser must be a function of the observations of the diagnoser. The following is an example of a noncausal design.

Example: There is only one agent. The first FSM in figure 2 is the plant model. There are two unobservable events, $\Sigma_{puo} = \{f_1, f_2\}$. Both are failure events. All other events are observable. We pick $\Sigma_{mf} = \{mf1, mf2\}$. The second FSM models a design (L). L is correct but noncausal since for two different plant behaviors, generating the same observation, it generates two different message sequences.

2.2 Correct Designs

It is desirable to design diagnosers that generate failure messages when there are failures in the plant and do not generate failure messages when there are no failures in the plant. These ideas are formalized as follows. It is assumed in the following that L is admissible.

Definition 2 *L is i-detecting iff for all $\sigma_{fi} \in \Sigma_{fi}$, there exists n_i such that for all $u, v \in \Sigma^*, u\sigma_{fi}v \in L, |P_{\Sigma_p}(v)| \geq n_i$, $\sigma_{mfi} \in u\sigma_{fi}v$ or there exists $t \in (\Sigma_m \cup \Sigma_{mf})^*$ such that $u\sigma_{fi}vt\sigma_{mfi} \in L$.*

The definition says that for all failures and for all ways in which that failures might occur, the design is such that it will generate a failure message within finitely many plant events following the failure, though it may wait to receive finitely many messages caused by the plant behavior before generating the failure message.

Definition 3 *L is i-false alarm free iff for all $u\sigma_{mfi} \in L, \sigma_{fi} \in u$.*

The definition says that the design is such that every generation of a fault message by a diagnoser is preceeded by the occurrence of the corresponding fault in the plant.

Definition 4 *L is ai-correct iff L is i-detecting and i-false alarm free.*

It is desirable that all diagnosers be correct. In our experience, a diagnostic design is generally correct under assumptions restricting the failures that can occur in the plant. Nevertheless it is important to understand the existence and computation of such solutions.

Note the similarities to the Neymann-Pearson problem [8]. This problem is deterministic. Therefore unlike the Neymann-Pearson problem languages with type one or type two errors cannot be compared in terms of false alarm and misdetection probabilities. Nevertheless in practise detecting languages with false alarms are a useful concept.

Like the Neyman-Pearson problem the following are true. A detecting language exists for every L_p, i.e., choose,

$$L = \{u : u \in L_p, \text{ or } u = v\sigma_1 \ldots \sigma_k, v \in L_p\},$$

$\Sigma_{mfi} = \{\sigma_1, \ldots \sigma_k\}$. Similarly, a false alarm free L always exists, i.e.,choose $L = L_p$.

2.3 Diagnosability

We next formalize independent diagnosability, decentralized diagnosability and centralized diagnosability. We begin with the definition of decentralized diagnosability since it is the most complex.

Definition 5 *The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is decentrally diagnosable iff there exists $\{\Sigma_{mi}\}_{i\in\mathcal{A}}$ and $L \subseteq \Sigma^*$ such that L is $i-$correct for all $i \in \mathcal{A}$.*

Thus a plant is decentrally diagnosable iff there exists some message set and associated design in which each site can correctly generate the failure messages of interest after waiting for finitely many plant events and finitely many inter-diagnoser message events caused by the finitely many plant events.

Definition 6 *The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is independently diagnosable iff for $\Sigma_{mi} = \emptyset$ for all $i \in \mathcal{A}$ there exists $L \subseteq \Sigma^*$ such that L is $i-$correct for all $i \in \mathcal{A}$.*

Thus a plant is indenpendently diagnosable if it is decentrally diagnosable with an empty inter-diagnoser communication message set, i.e., without inter-diagnoser communication.

Definition 7 *The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is centrally diagnosable iff the corresponding single site triple $(L_p, \Sigma_{po}, \Sigma_f)$ is decentrally diagnosable, i.e., there exists L with $\Sigma_m = \emptyset$ such that L is $1-$correct.*

Thus a plant is centrally diagnosable iff a central diagnoser that observes every event observed by every site instantaneously, i.e., it observes the exact ordering of the observations of different sites, can diagnose all the failures of interest to all the sites.

It is immediate from definitions 6 and 5 that independent diagnosability implies decentralized diagnosability. The fact that decentralized diagnosability implies centralized diagnosability is also almost immediate from the definitions. It may be argued as follows. Let $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ be decentrally diagnosable and $L, \{\Sigma_{mi}\}_{i\in\mathcal{A}}$ be the corresponding design. Let $\Sigma' = \Sigma_p \cup \Sigma_{mf}$. Then $P_{\Sigma'}(L)$ is $1 - correct$ for the triple $(L_p, \Sigma_{po}, \Sigma_f)$. In other words the centralized diagnostic design is obtained by simply deleting all the inter-diagnoser messaging events. The messaging is redundant because all the plant events causing the messages were centrally observed anyway.

3 Motivating Example

This section discusses an automated vehicle platoon wireless LAN diagnostic problem [3] that motivates our formulation of the diagnosis problem for distributed systems.

The concept of automated vehicle platoons on intelligent highways has been a subject of aclantive research for several years. The following linear control law [14] has been extensively used for platoon follower control,

$$u_i = (1-k_a)\ddot{x}_l + k_a\ddot{x}_{i-1} + k_v(\dot{x}_{i-1}-\dot{x}_i) + k_p(x_{i-1}-x_i-L)$$
$$+c_v(\dot{x}_l-\dot{x}_i) + c_p(x_l-x_i-iL),$$

where k_a, k_v, k_p, c_v, c_p are control gains, x_l is the position of the lead vehicle in an inertial reference frame, and x_i is the position of the $i-$th platoon follower in the same reference frame. Observe that the commanded acceleration of a platoon follower vehicle is a function of variables that cannot be sensed, such as the acceleration and velocity of the lead vehicle. Therefore platoon operation is supported by a wireless LAN [4],[7] that transmits information from one vehicle to the other at the sampling rate (typically 20ms) of the longitudinal control system. We are interested in the problem of designing diagnostics for the wireless LAN.

The physical structure of the platoon LAN corresponds closely with the distributed system diagnostic architecture in figure 1. Each vehicle in the platoon constitutes a site. It has radios that can transmit and receive. Each vehicle needs to have a diagnoser that can diagnose faults in its radios. The site diagnoser observes only the messages received by its radio. We shall see that the LAN radio network is centrally diagnosable but not independently diagnosable. A central diagnoser is difficult to realize and undesirable in a distributed system that is required to be highly fault tolerant. Therefore a decentralized diagnoser design is required and the vehicle diagnosers must exchange information for the purpose of diagnosis. We are motivated to understand if for a system with a global observation structure rich enough to support centralized diagnosis there always exists an inter-diagnoser messaging scheme that can replicate some of the information utilized to obtain central diagnosability, i.e., are centrally diagnosable systems decentrally diagnosable ? We explain the intution behind our approach to this question in the context of the LAN example. Section 4 provides a more formal answer to this question.

We use a simplified model of the LAN operation that captures the features essential for diagnostic design. For the full model and design see [3]. Using this model, we show that the LAN diagnosis problem can be solved if the diagnosers can exchange information over a network distinct from the LAN itself. Fortunately the vehicles are also connected to a WAN that can be used for occaisional messaging. It should be noted that there is no dedicated WAN bandwidth for a vehicle and therefore there are no deterministic bounds on the WAN communication delays.

sync	trans. of synchronization pulse by lead
f1:r:s	recep. of synchronization by first follower
f2:r:s	recep. of synchronization by second follower
t	expiration of the time slot
l:m	trans. of the lead vehicle control message
f1:r:l	recep. of lead vehicle message by first
f2:r:l	recep. of lead vehicle message by second
f1:m	trans. of control message by first follower
l:r:f1	recep. of first follower message by lead
f2:r:f1	recep. of first follower message by second
f2:m	trans. of control message by second follower
l:r:f2	recep. of second follower message by lead
f1:r:f2	recep. of second follower message by first

Table 1: LAN Model Events

The operation of the LAN network for a three vehicle platoon is shown by the finite state machine (FSM) in figure 3. The meanings of the event labels are explained by table 1. The vehicles share the LAN through a TDMA scheme [4, 7]. The longitudinal control sampling interval is divided into slots of fixed duration with one slot being allocated to each vehicle in the platoon. At the beginning of the sampling interval the lead vehicle transmits a synchronizing pulse to which all the platoon vehicles set their clocks. The first slot is then used by the lead vehicle to transmit control information, the second is used by the first follower vehicle of the platoon, and so on until the last vehicle transmits, after which the lead vehicle synchronizes clocks again. Therefore under normal conditions the LAN has a time-driven operation with vehicles consecutively transmitting on the shared LAN channel. Since all transmissions occur on a shared channel we assume that in the abcense of faults every transmission by a vehicle is received by every other vehicle.

We will illustrate the diagnosis of faults in the lead and first follower vehicles. The diagnosis problem for subsequent followers is similar to that of the first follower. It is assumed that each radio can have a transmitter or receiver fault. The former implies that it is unable to transmit and the latter that it is unable to receive. Therefore we are concerned with diagnosis of the four fault events in table 2.

ltf	lead vehicle transmitter fault
lrf	lead vehicle receiver fault
f1tf	first follower transmitter fault
f1rf	first follower receiver fault

Table 2: LAN Failure Events

A vehicle that does not receive the synchronising pulse does not transmit since it may collide with the others if it does so. The lead vehicle transmits even if it does not receive anything because it controls the synchronisation. Once again for simplicity, it is assumed here that the faults occur only at the beginning of the

sampling interval and only single failures are considered. If multiple failures, occur they will be diagnosed incorrectly. Fortunately multiple failures are rare. In our experience it is generally the case that most systems cannot be instrumented to isolate multiple simultaneous failures.

Since each vehicle has a receiver and a transmitter, it is assumed that each vehicle is desirous of diagnosing its own receiver and transmitter fault. Moreover there is a natural decomposition of observation, i.e., each vehicle diagnoser only observes the messages received by its receiver. Transmission events are unobservable. The clock ticks are observable. Formally the diagnostic problem is specified by the plant FSM of figure 3, and the following sets of failure and observable events for each vehicle

$$\begin{aligned}\Sigma_{ol} &= \{tick, l:r:f1, l:r:f2\},\\ \Sigma_{of1} &= \{tick, f1:r:l, f1:r:f2\},\\ \Sigma_{of2} &= \{tick, f2:r:f1, f1:r:f2\},\\ \Sigma_{fl} &= \{ltf, lrf\}, \Sigma_{ff1} = \{f1tf, f1rf\},\\ \Sigma_{ff2} &= \{f2tf, f2rf\}.\end{aligned}$$

The projection of G_{lan} onto Σ_{ol} indicates that over an entire sampling interval the lead vehicle observes the same behavior for the failures ltf and lrf. Likewise it observes the same behavior for the failures $f1rf$ and $f1tf$. To reduce complexity we encode the observation process over the entire sampling interval into single events. Since there are three distinct kinds of observation over a sampling interval for the lead vehicle, the encoding is represented by three events a, b, c. In the case of the first follower the failures ltf and $f1rf$ are indistinguishable, as are the normal mode and the failures $f1tf, lrf$. Therefore for the first follower the encoding of the observation process can be represented by two events d, e. We study the diagnostic problem by projecting the plant onto the reduced event set $\{ltf, lrf, f1tf, f1rf, a, b, c, d, e\}$. The reduced plant FSM (G^{red}_{lan}) is shown in figure 5. Figure 4 represents the plant property that in every sampling interval observations will be made by both lead and follower vehicles. This sort of "fairness" in observation is important since if the plant has arbitrarily long behaviors that only yield observations to a single diagnoser, then the other diagnosers are of little use in these situations. The local information of the single diagnoser must suffice.

The LAN is not independently diagnosable. The lead vehicle observes the event c for both ltf and lrf, making them indistinguishable given local information. Similarly, the first follower observes e for the normal mode, $f1tf$, and lrf making these two failures indistinguishable from each other and the normal mode given local information. On the other hand the LAN is centrally diagnosable since a diagnoser observing all the messages received by the lead and follower vehicles, i.e., with the observable event set $\Sigma_{ol} \cup \Sigma_{of1}$, could distinguish all four failure modes and the normal mode. Figure 5 shows that the future possible observable event sequences from the state sets

$$\{0,9\}, \{1,5\}, \{2,6\}, \{3,7\}, \{4,8\},$$

are all pairwise different. For any of the four failures, the third event following the failure will provide a signature that will uniquely identify the state entered by the plant after the failure and therefore the failure event itself. In section 4 we prove the diagnosability of this system formally.

Finally, we are left with the question of whether the LAN is decentrally diagnosable, i.e., whether there exists an inter-diagnoser messaging scheme that can provide each vehicle with some of the information that enables central diagnosability. One straightforward communication scheme is for each vehicle to communicate all its observations to the others. Intuitively, this scheme represents the possibility of "maximizing" the replication of global information. We show in section 4 that if this communication scheme does not work then no communication scheme works. However, this communication scheme is not a practical solution to decentralized diagnosis problems for at least two reasons. Firstly, the bandwidth required to support diagnosis would likely be very nearly the same as that required for normal operation. The second reason arises explicitly from our modeling assumptions. Suppose the lead vehicle observes the sequence b^n, and is designed to communicate every observation of b to the first follower by generating the message event m_b. Since we make no assumptions on the delays in the inter-diagnoser communication network, there is clearly no bound on the number of $b's$ that may be generated by the plant before the first m_b is delivered to the other diagnoser. Since each b observation puts an m_b in the message queue, the lead vehicle diagnoser in general needs to schedule an unbounded queue (specifically of size n for b^n) of m_b messages to be delivered. Formally, we show in section 4 that for a regular plant language L_p, the correct L corresponding to the "communicate all observations" communication scheme is not generally regular.

A more practical decentralized diagnosis scheme that works for this example is shown in figures 6 and 7. The two figures show diagnoser designs for the lead and follower vehicle respectively. Basically, if the lead observes c it sends a message to the first follower. The follower then reports the first observation after receipt of the message. The lead vehicle then generates the appropriate fault message. If the lead vehicle observes a then it reports the observation to the first follower. The first follower then waits for the first plant observation following the lead vehicles message and then generates the appropriate failure message. The protocol is similar if the lead observes c.

4 Existence of a Correct Design

The following theorem provides a necessary and sufficient condition for the existence of a messaging scheme

for a given decentralized diagnostic problem. The condition is a qualitative property of the DES plant. It generalizes a theorem presented in [12] on centralized diagnosis. We apply it to the example in section 3 and to other examples that help us understand the distinction between systems for which inter-diagnoser communication can replicate some of the global information adequately and systems for which it cannot. We also state a corollary of the theorem stating necessary and sufficient conditions for the three kinds of diagnosability.

The theorem states that a plant is decentrally diagnosable iff for every failure event, and every plant behavior that preceeds and succeeds the failure event, where the succeeding behavior is sufficiently long, any other behavior that looks the same to all the sites must also have the same failure in it. Thus as long as every failed and non-failed plant behavior can be distinguished by at least one site the diagnosis problem can be correctly solved. In particular, we show the sufficiency of the plant property by proving that the "communicate all observations" messaging scheme results in a correct diagnostic design.

Theorem 1 *There exists L, $i-$correct for all $i \in \mathcal{A}$, iff there exists $n \in \mathbb{N}$ such that for all $\sigma_f \in \Sigma_f$ and $u\sigma_f v \in L_p$ with $|v| > n$, the following condition is satisfied,*

$$(w \in L_p) \wedge (\forall i, P_{\Sigma_{poi}}(u\sigma_f v) = P_{\Sigma_{poi}}(w)) \Rightarrow \sigma_f \in w.$$

The proof is omitted. Sufficiency is proved by constructing a canonical "communicate all observations" solution to the decentralized diagnosis problem. This is in a sense a maximal communication solution. The construction is as follows. Pick a collection of message sets $\{\Sigma_{mi}\}_{i\in\mathcal{A}}$ such that the message sets are disjoint from each other, Σ_p, and Σ_{mf}. Furthermore each Σ_{mi} should be such that there exists a bijection $\eta_i : \Sigma_{poi} \to \Sigma_{mi}$. Let

$$\begin{aligned} L = \{w \in \Sigma^* : & P_{\Sigma_p}(w) \in L_p, \forall i P_{\Sigma_{mi}}(w) \leq \eta_i \circ P_{\Sigma_{poi}}(w), \\ & \text{and } (w = s\sigma_{mf}) \Rightarrow \exists u\sigma_f v \in L_p, |v| > n, \\ & \text{s.t. } \forall j((P_{\Sigma_{poj}}(s) \geq P_{\Sigma_{poj}}(u\sigma_f v)) \\ & \wedge (P_{\Sigma_{mj}}(w) \geq \eta_j \circ P_{\Sigma_{poj}}(u\sigma_f v))\}. \end{aligned}$$

This language is not regular in general.

The following corollary of theorem 1 relates the theorem to the definitions of decentralized, centralized, and independent diagnosability.

Corollary 1 *1. The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is decentrally diagnosable iff there exists $n \in \mathbb{N}$ such that for all $\sigma_f \in \Sigma_f$ and $u\sigma_f v \in L_p$ with $|v| > n$, the following condition is satisfied,*

$$(w \in L_p) \wedge (\forall i, P_{\Sigma_{poi}}(u\sigma_f v) = P_{\Sigma_{poi}}(w)) \Rightarrow \sigma_f \in w.$$

2. The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is centrally diagnosable iff there exists $n \in \mathbb{N}$ such that for all $\sigma_f \in \Sigma_f$ and $u\sigma_f v \in L_p$ with $|v| > n$, the following condition is satisfied,

$$(w \in L_p) \wedge (P_{\Sigma_{po}}(u\sigma_f v) = P_{\Sigma_{po}}(w)) \Rightarrow \sigma_f \in w.$$

3. The triple $(L_p, \{\Sigma_{poi}\}_{i\in\mathcal{A}}, \{\Sigma_{fi}\}_{i\in\mathcal{A}})$ is independently diagnosable iff for all $i \in \mathcal{A}$ there exists $n_i \in \mathbb{N}$ such that for all $\sigma_f \in \Sigma_{fi}$ and $u\sigma_f v \in L_p$ with $|v| > n_i$, the following condition is satisfied,

$$(w \in L_p) \wedge (P_{\Sigma_{poi}}(u\sigma_f v) = P_{\Sigma_{poi}}(w)) \Rightarrow \sigma_f \in w.$$

The following is an example of a plant that can be diagnosed by a single diagnoser but cannot be diagnosed by two diagnosers who jointly observe the same events as the single agent but observe less severally.

Example: There are two diagnosers with observation and failure events events $\Sigma_{po1} = \{a, c\}, \Sigma_{f1} = \{f1\}, \Sigma_{po2} = \{b, c\}, \Sigma_{f2} = \{f2\}$. The plant is shown in figure 8. There is no correct messaging design as evident from $u\sigma_f v = f1abc^n$, and $w = f2bac^n$. Then for all n, $P_{\Sigma_{oi}}(u\sigma_f v) = P_{\Sigma_{oi}}(w)$, $i = 1, 2$ and $f1 \notin w$. However, if there is a single diagnoser with observation $\Sigma_{po} = \{a, b, c\}$, a correct design exists because $P_{\Sigma_{po}}(u\sigma_f v) \neq P_{\Sigma_{po}}(w)$ for any $n \in \mathbb{N}$.

This is an example of plant in which the isolation of $f1$ and $f2$ from each other depends upon the correct ordering of the observable events generated by the plant and with the decentralization of observation being as given no amount of messaging or memory can reconstruct the exact order in which a and b occurred. It seems that the only way to diagnose failures in this kind of plant would be to use time-stamps from a global clock.

We can use theorem 1 to show that a correct diagnostic design exists for the LAN network example discussed in section 3. The argument refers to the finite state machines shown in figure 5. We pick the $n = 2$, where n is as in theorem 1. Consider $\sigma_f = f1tf$. Then for all $uf1tf \in L_p$, $v = ae$ or $v = ea$. Let w be such that

$$P_{\Sigma_{ol}}(u\sigma_f v) = P_{\Sigma_{ol}}(w), P_{\Sigma_{f1}}(u\sigma_f v) = P_{\Sigma_{f1}}(w), \sigma_f \notin w.$$

Since $a \in w$, the assumptions imply $f1rf \in w$. Since $d \notin w$, this implies

$$w = (eb)^n f1rf, \text{ or } w = (eb)^n ef1rf.$$

This implies $a \notin w$ which is a contradiction. Therefore $f1tf \in w$. The arguments are similar for the three other failure events. They are not presented.

5 Summary

We have formulated a diagnostic problem for distributed systems within the context of a language-theoretic discrete event formalism. The diagnosis problem is non-trivial because the plant is partially observed.

In particular the failure events are assumed to be unobservable. We say that the distributed system is diagnosable if there exists a failure message design that is a function of the observations that is detecting and false-alarm-free in the Neymann-Pearson sense. Three notions of diagnosability, namely centralized, decentralized, and independent, are investigated.

In a plant with decentralized observation the sites collectively observe more than they do individually. We have presented a wireless LAN diagnosis problem in which local site information is inadequate to diagnose the failures at the site, but the collective information is adequate. However, the collective information can only be realized by inter-diagnoser messaging. A suitable message design is presented for the wireless LAN example. We show that in general there exist finite state systems for which the full collective information cannot be realized by any inter-diagnoser communication scheme. In other words, there exist centrally diagnosable systems that are not decentrally diagnosable. This happens when diagnosis depends on the ordering of plant events and the order cannot be reconstructed due to the decentralization of information. This problem could not arise in a distributed system where the site clocks are synchronized. However, in systems with unsynchronized local clocks these problems are to be expected.

We have presented a theorem that qualitatively describes, in a necessary and sufficient manner, the class of partially observed discrete event plants for which there exists an inter-diagnoser messaging scheme that is adequate. This result generalizes the results presented in [12]. We are able to use the theorem to prove the diagnosabiliity or non-diagnosability of some interesting examples that give us an understanding of the distinctions between the solvable and unsolvable problems. We investigated the properties of the "communicate all observations scheme" that intutively seems to maximize the information of each site. The proof of theorem 1 shows that in a decentrally diagnosable plant this scheme always works though it may require unbounded memory to execute. For the LAN example we can see that better finite memory solutions exist.

The investigations described in this paper offer many interesting avenues of further research. Within the context of this model one could investigate decidability and the synthesis of efficient communication. We are also left with the feeling that an interleaved semantics is a rather poor way of representing concurrency. Richer models expressing bounds on communication delays or concurent observation would allow more tractable formulations of partial observation problems in distributed systems.

References

[1] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, and C. Jard. A petri net approach to fault detection and diagnosis in distributed systems (part 2). *Proc. 36th IEEE CDC*, December 1997.

[2] R. Boubour, C. Jard, A. Aghasaryan, E. Fabre, and A. Benveniste. A Petri net approach to fault detection and diagnosis in distributed systems (part 1). In *Proc. 36th IEEE CDC*, December 1997.

[3] F. Eskafi. A diagnostic system design for the intra-platoon communication system in NAHSC demo'97. Preprint, PATH, UCBerkeley, December 1997.

[4] B. Foreman. A survey of wireless communication technologies for automated vehicle control. In *Proc. SAE FTT*, August 1995.

[5] D. N. Godbole and R. Sengupta. Tools for the design of fault management systems. In *Proc. 1997 IEEE ITSC*, November 1997.

[6] A. E. Lindsey and P. Vishwanath. Design, verification and failure diagnosis of wireless communication protocols for ahs. In *Proc. 1997 IEEE ITSC*, November 1997.

[7] PATH wireless vehicle communication system: Overview and functional specifications, MPI Corp, Atlanta, GA. October 1996.

[8] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. Roy. Soc., Ser. A.*, 231:289–337, 1933.

[9] D. N. Pandalai and L. E. Holloway. Template languages for fault monitoring of single-instance and multiple-instance discrete event processes. In *Proc. 36th IEEE CDC*, December 1997.

[10] R. Rajamani, J.K. Hedrick, and A. Howell. A complete fault diagnostic system for longitudinal control of automated vehicles. In *Proc. Symposium of Advanced Automotive Technologies, 1997 ASME International Congress*, 1997.

[11] P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proc. IEEE*, 77(1), Jan 1989.

[12] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE TAC*, 40, September 1995.

[13] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Failure diagnosis using discrete-event models. *IEEE TCST*, 4, March 1996.

[14] D. Swaroop. *String stability of interconnected systems.* PhD thesis, UCBerkeley, December 1994.

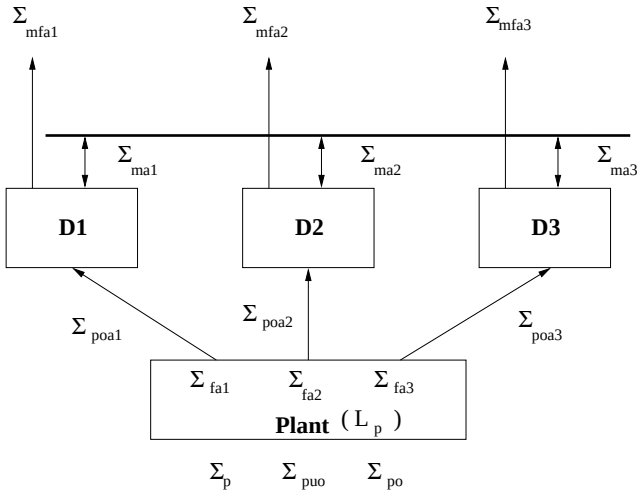

Figure 1: Decentralized Diagnostic System Architecture

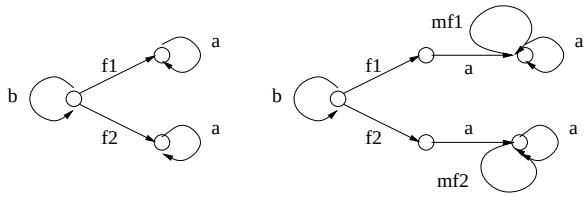

Figure 2: A plant and a non-causal diagnoser design

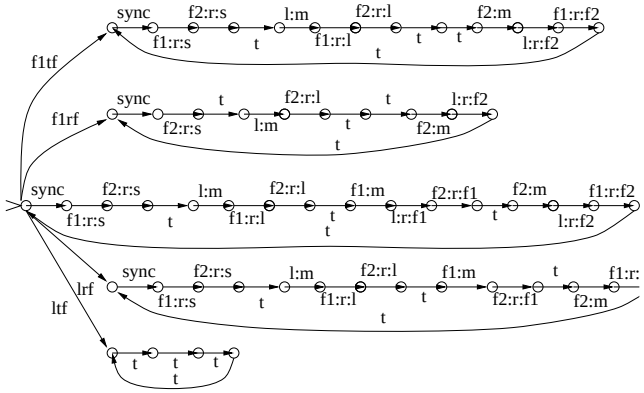

Figure 3: Complete LAN Model for a Three Vehicle Platoon

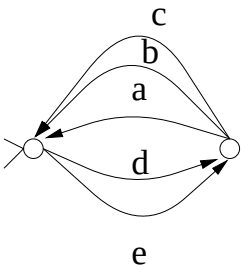

Figure 4: Projection of the Plant onto the Reduced Observable Event Set

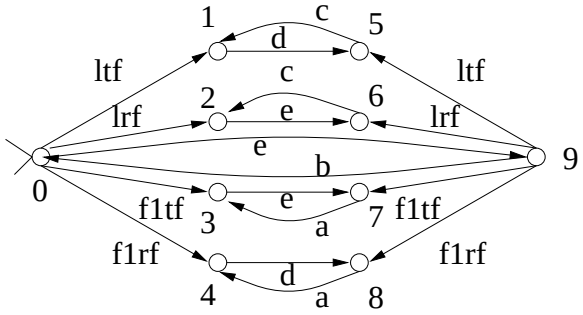

Figure 5: Reduced LAN Model

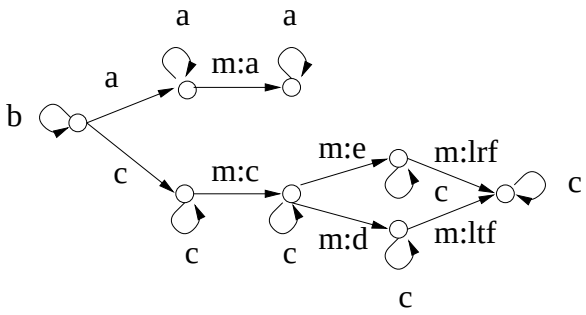

Figure 6: Lead Vehicle LAN Diagnoser Design

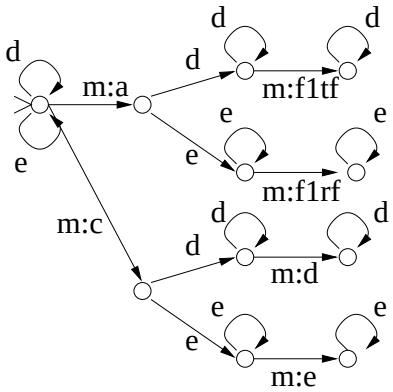

Figure 7: First Follower Vehicle LAN Diagnoser Design

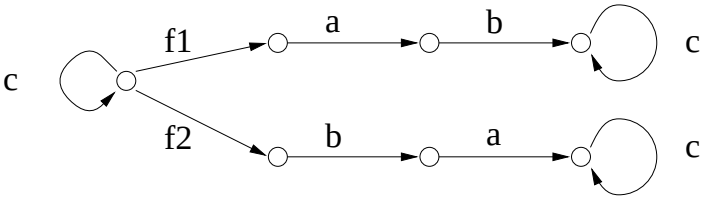

Figure 8: The difference between centralized and decentralized diagnostics