

Decentralized diagnosis of Petri nets

Maria Paola Cabasino, Alessandro Giua, Andrea Paoli, Carla Seatzu

Abstract

In this paper we deal with the problem of failure diagnosis of discrete event systems with decentralized information. The decentralized architecture that we use is composed by a set of sites communicating their diagnosis information with a coordinator that is responsible of detecting the occurrence of failures in the system. In particular, we define two protocols that differ for the amount of information exchanged between the local sites and the coordinator, and the rules adopted by the coordinator to compute the global diagnosis states.

Published as:

M.P. Cabasino, A. Giua, A. Paoli, C. Seatzu "Decentralized diagnosis of Petri nets," ACC10: 2010 American Control Conference (Baltimore, MD, USA), Jun-Jul 2010. To appear.

This work has been partially supported by the European Community's Seventh Framework Programme under project DISC (Grant Agreement n. INFSO-ICT-224498).

M.P. Cabasino, A. Giua and C. Seatzu and are with the Department of Electrical and Electronic Engineering, University of Cagliari, Piazza D'Armi, 09123 Cagliari, Italy {seatzu,cabasino@diee.unica.it}.

A. Paoli is with the Center for Research on Complex Automated Systems - Department of Electronic, Computer Science and Systems, University of Bologna, Via Risorgimento, 2, 40136 Bologna, Italy {andrea.paoli@unibo.it}.

I. INTRODUCTION

The problem of failure detection has received a lot of attention in industrial systems in the past few decades. Solving a problem of diagnosis means that we associate to each observed string of events a diagnosis state, such as “normal” or “faulty” or “uncertain”. In the literature a lot of contributions have been presented for discrete event systems in the centralized framework, e.g., [1]–[5]. Due to the intrinsic distributed nature of the real systems, a lot of distributed diagnosis techniques, that take advantage of the natural decompositions of a modular system, have been studied both dealing with automata [6]–[10] and Petri nets [11]–[13].

In this paper we present an approach for diagnosis of Petri nets with decentralized information that combines the work of Debouk *et al.* [8] with the approach presented by some of us in [2], [14]. In particular, we start from the same decentralized architecture considered in [8] and from a series of similar assumptions on the considered model. However, here we solve the decentralized diagnosis problem in the context of Petri nets, while the approach in [8] is in the framework of automata. This enables us to keep the advantages of the centralized approach we proposed in [2], [14].

We assume that the system is monitored by a set of sites. Each site knows the structure of the net and the initial marking but observes the evolution of the system with a different mask, i.e., the set of observable transitions is different for each site. Diagnosis is locally performed using the approach we previously introduced in [2], [14] whose main feature is that of avoiding an exhaustive enumeration of the set of sequences that may have fired given the actual observation. It is also based on the definition of four diagnosis states, each of which can be associated with a number from 0 to 3, depending on the degree of alarm. For instance, 3 is used to capture the fact that the fault has occurred for sure, whereas 0 captures the fact that the fault has not occurred for sure.

Using its own observation, each site performs diagnosis and, according to a given protocol, communicates it, eventually with some other information, to the coordinator who calculates global diagnosis states. In particular, two different protocols are defined that differ for the amount of information exchanged between the coordinator and the local sites, and viceversa. In both cases an important property is proved, namely that the coordinator never produces false alarms. Finally, the diagnosability property under decentralization is investigated.

II. BACKGROUND ON LABELED PETRI NETS

A *Place/Transition net* (P/T net) is a structure $N = (P, T, Pre, Post)$, where P is the set of m places, T is the set of n transitions, $Pre : P \times T \rightarrow \mathbb{N}$ and $Post : P \times T \rightarrow \mathbb{N}$ are the pre and post incidence functions that specify the arcs. The function $C = Post - Pre$ is called incidence matrix.

A *marking* is a vector $M : P \rightarrow \mathbb{N}$ that assigns to each place a nonnegative integer number of tokens; the marking of a place p is denoted with $M(p)$. A *net system* $\langle N, M_0 \rangle$ is a net N with initial marking M_0 .

A transition t is enabled at M iff $M \geq Pre(\cdot, t)$ and may fire yielding the marking $M' = M + C(\cdot, t)$. The notation $M[\sigma\rangle$ is used to denote that the sequence of transitions $\sigma = t_1 \dots t_k$ is enabled at M ; moreover we write $M[\sigma\rangle M'$ to denote the fact that the firing of σ from M yields to M' .

The set of all sequences that are enabled at the initial marking M_0 is denoted with $L(N, M_0)$. Given a sequence $\sigma \in T^*$, we call $\pi : T^* \rightarrow \mathbb{N}^n$ the function that associates to σ a vector $y \in \mathbb{N}^n$, named *firing vector*, such that $y(t) = k$ if the transition t is contained k times in σ .

A marking M is said to be *reachable* in $\langle N, M_0 \rangle$ iff there exists a firing sequence σ such that $M_0[\sigma\rangle M$. The set of all markings reachable from M_0 defines the *reachability set* of $\langle N, M_0 \rangle$ and is denoted with $R(N, M_0)$. Finally we define $PR(N, M_0)$ the potentially reachable set, i.e., the set of all markings $M \in \mathbb{N}^m$ for which there exists a vector $y \in \mathbb{N}^n$ that satisfies the *state equation* $M = M_0 + C \cdot y$. It holds that $R(N, M_0) \subseteq PR(N, M_0)$.

A PN having directed circuits is called *acyclic*. For such nets if the vector $y \in \mathbb{N}^n$ satisfies the equation $M_0 + C \cdot y \geq 0$, there exists a firing sequence σ firable from M_0 and such that the firing vector associated with σ is equal to y . Moreover for acyclic nets $R(N, M_0) = PR(N, M_0)$.

A *labeling function* $\mathcal{L} : T \rightarrow L \cup \{\varepsilon\}$ assigns to each transition a symbol from a given alphabet L or the empty string ε . The set of transitions sharing the same label l is denoted as T_l . Transitions whose label is ε are called *silent* and are denoted by the set T_u . The set $T_o = T \setminus T_u$ is the set of *observable transitions*, i.e., when an observable transition fires we observe its label. We denote as C_u (C_o) the restriction of the incidence matrix to T_u (T_o). Moreover, given a sequence $\sigma \in T^*$, $P_u(\sigma)$ ($P_o(\sigma)$) denotes the projection of σ over T_u (T_o).




Fig. 1. The decentralized diagnosis architecture.

We denote as $w = \mathcal{L}(\sigma)$ the word of events associated to the sequence σ . We define

$$\mathcal{S}(w) = \{\sigma \in L(N, M_0) \mid \mathcal{L}(\sigma) = w\}$$

the set of sequences consistent with $w \in L^*$. In plain words, given an observation w , $\mathcal{S}(w)$ is the set of sequences that may have fired.

Finally, given a net $N = (P, T, Pre, Post)$ and a subset $T' \subseteq T$ of its transitions, we define the T' -induced subnet of N as the new net $N' = (P, T', Pre', Post')$, where Pre' and $Post'$ are the restrictions of Pre and $Post$ to T' , i.e., N' is the net obtained from N removing all transitions in $T \setminus T'$. We write that $N' \prec_{T'} N$.

III. PROBLEM STATEMENT

We model anomalous or faulty behavior using the set of silent transitions $T_f \subseteq T_u$. The set T_f includes all fault transitions and is further decomposed into r different subsets T_f^i , where $i \in \mathcal{F} = \{1, \dots, r\}$, that model different fault classes. The transition set $T_{reg} = T_u \setminus T_f$ represents the set of unobservable, but regular, transitions.

The problem of fault diagnosis can be seen as the problem of detecting the firing of any fault transition in T_f , using the knowledge on the firing of observable transitions, or the knowledge on their labels in the case of labeled Petri nets.

In this work we explore the possibility of performing diagnosis using a decentralized architecture as depicted in Fig. 1. The system is monitored by a set $\mathcal{J} = \{1, \dots, \nu\}$ of sites. Each

site has a complete knowledge of the net structure and of the initial marking, but observes the evolution of the system using its own observation mask. Obviously, different sites have different observation masks. In particular, for any site $j \in \mathcal{J}$, the set of locally observable transitions is the set $T_{o,j} \subseteq T_o$. Any centrally observable transition is observed by at least one site, i.e., $\bigcup_{j \in \mathcal{J}} T_{o,j} = T_o$. The set of locally unobservable transitions is defined as

$$T_{u,j} = T_{reg} \cup T_f \cup (T_o \setminus T_{o,j}). \quad (1)$$

We denote as $L_j \subseteq L$ ($j \in \mathcal{J}$) the alphabet of the j -th site, i.e., the set of labels observable by the j -th site. Moreover, we denote as $w_j = \mathcal{L}_j(\sigma)$ the word of events in L_j associated to the sequence σ by the j -th site.

As shown in Fig. 1, on the basis of its own observation $w_j = \mathcal{L}_j(\sigma)$ ($j \in \mathcal{J}$) each site performs a local diagnosis. In particular, for each fault class $i \in \mathcal{F}$ it computes a different diagnosis state $\Delta_{j,i}$ and depending on this, it exchanges information with a *coordinator* C according to a given protocol¹. The coordinator fuses the information coming from the different sites according to the considered protocol and infers on the occurrence of faults. More precisely, for each fault class $i \in \mathcal{F}$ it computes a diagnosis state $\bar{\Delta}_i$.

In this paper we explore the decentralized architecture under the following assumptions.

- A1** The same label $l \in L$ can be associated to more than one transition, but if a site observes a transition labeled l , then it observes any transition whose label is l , namely, $\nexists t, t'$ such that $\mathcal{L}(t) = \mathcal{L}(t')$ and $t \in T_{o,j}$, while $t' \notin T_{o,j}$.
- A2** The $T_{u,j}$ -induced subnet $N_{u,j}$ is acyclic for any $j \in \mathcal{J}$.
- A3** The coordinator C knows which transitions can be observed by each site, i.e., it knows the sets $T_{o,j}$ for any $j \in \mathcal{J}$.
- A4** There is reliable communication between the local sites and the coordinator, i.e., all messages sent from a local site are received by the coordinator, and viceversa, correctly and in order.

Note that we also investigate an important issue that occurs when performing diagnosis, regardless of the fact that it is centralized or decentralized, namely that of *diagnosability*.

¹For the sake of simplicity in Fig. 1 we represented the diagnosis states in a vectorial form, thus $\Delta_{j,i}$ denotes the i th component of Δ_j . The same notation has been used for the diagnosis state computed by the Coordinator C .

Definition 3.1: Let us consider a Petri net system $\langle N, M_0 \rangle$ having no deadlock after the occurrence of transition $t_f \in T_f^i$, for all $i \in \mathcal{F}$. Assume that diagnosis is performed according to a given approach (either centralized or decentralized).

We say that $\langle N, M_0 \rangle$ is *diagnosable with respect to* (wrt) *the fault class* T_f^i *and wrt a given diagnosis approach* iff the occurrence of some fault in T_f^i is unambiguously detected using the specified diagnosis approach after a *finite* number of transition firings. ■

Definition 3.2: A Petri net system $\langle N, M_0 \rangle$ is *diagnosable wrt a given diagnosis approach* if it is diagnosable wrt that approach for all fault classes T_f^i , $i \in \mathcal{F}$. ■

Note that in the centralized framework, inspired by the definition of diagnosability for languages introduced in [15], Definition 3.1 can alternatively be formulated as follows.

Definition 3.3: A Petri net system $\langle N, M_0 \rangle$ having no deadlock after the occurrence of transition $t_f \in T_f^i$, for $i \in \mathcal{F}$, is *diagnosable wrt the fault class* T_f^i if there do not exist two firing sequences σ_1 and $\sigma_2 \in T^*$ satisfying the following conditions:

- $\mathcal{L}(\sigma_1) = \mathcal{L}(\sigma_2)$,
- $\forall t_f \in T_f^i, \sigma_1 \in (T \setminus T_f^i)^*$,
- \exists at least one $t_f \in T_f^i$ such that $t_f \in \sigma_2$,
- σ_2 is of “arbitrary length” (see [15]) after fault $t_f \in T_f^i$.

■

IV. BASIC DEFINITIONS AND RESULTS ON CENTRALIZED DIAGNOSIS

In this section we briefly recall the diagnosis procedure we defined in [2], [14] in the centralized framework, that is used by the different sites to perform diagnosis locally. As in the previous section, $T = T_o \cup T_u$ where $T_u = T_{reg} \cup T_f$, and the observations coincide with the labels associated to transitions in T_o . In particular, we first provide some preliminary definitions.

- Given a word $w \in L^*$, let $\sigma_o \in T_o^*$ be a sequence of observable transitions such that $\mathcal{L}(\sigma_o) = w$. We call *justification of* w the sequence σ_u of unobservable transitions interleaved with σ_o whose firing enables σ_o and whose firing vector is minimal.

Since in general σ_o is not unique and more than one σ_u may be associated to each σ_o , then the set of justifications of w is not a singleton.

- We denote as $Y(M_0, w)$ the set of firing vectors relative to justifications of w .

The generic element $y \in Y(M_0, w)$ is called *j-vector*.




Fig. 2. The Petri net system considered in Examples 4.1 and 4.3.

- Finally, we denote as

$$\begin{aligned} \hat{\mathcal{J}}(w) = & \{ (\sigma_o, \sigma_u), \sigma_o \in T_o^*, \mathcal{L}(\sigma_o) = w, \\ & \sigma_u \in T_u^* \mid \\ & [\exists \sigma \in \mathcal{S}(w) : \sigma_o = P_o(\sigma), \sigma_u = P_u(\sigma)] \wedge \\ & [\nexists \sigma' \in \mathcal{S}(w) : \sigma_o = P_o(\sigma'), \sigma'_u = P_u(\sigma') \wedge \\ & \pi(\sigma'_u) \preceq \pi(\sigma_u)] \} \end{aligned}$$

the set of couples (sequence $\sigma_o \in T_o^*$ with $\mathcal{L}(\sigma_o) = w$ - corresponding *justification* of w).

Example 4.1: Let us consider the PN in Fig. 2, where the set of observable transitions is $T_o = \{t_1, t_2, t_3\}$ and the set of unobservable transitions is $T_u = \{\varepsilon_4, \varepsilon_5, \varepsilon_6, \varepsilon_7, \varepsilon_8\}$. The labeling function is $\mathcal{L}(t_1) = a$ and $\mathcal{L}(t_2) = \mathcal{L}(t_3) = b$.

Let $w = ab$ be the observed word. There exist two sequences that are consistent with the actual observation and whose firing vector is minimal, namely $\sigma' = \varepsilon_4 t_1 t_2$, $\sigma'' = \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3$. Thus $\sigma'_u = \varepsilon_4$ and $\sigma''_u = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8$ are the two justifications of w . The set of j-vectors is $Y_{min}(M_0, w) = \{[1 \ 0 \ 0 \ 0 \ 0]^T, [1 \ 0 \ 1 \ 1 \ 1]^T\}$, where $y' = [1 \ 0 \ 0 \ 0 \ 0]^T$ is relative to σ'_u , while $y'' = [1 \ 0 \ 1 \ 1 \ 1]^T$ is relative to σ''_u . Finally, $\hat{\mathcal{J}}(w) = \{(t_1 t_2, \varepsilon_4), (t_1 t_3, \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8)\}$. ■

Let us now recall the notions of *diagnoser* and *diagnosis states*.

Definition 4.2: A *diagnoser* is a function $\Delta : L^* \times \{T_f^1, T_f^2, \dots, T_f^r\} \rightarrow \{0, 1, 2, 3\}$ that associates to each observation w and to each fault class T_f^i , $i = 1, \dots, r$, a *diagnosis state*.

- $\Delta(w, T_f^i) = 0$ if for all $\sigma \in \mathcal{S}(w)$ and for all $t_f \in T_f^i$ it holds $t_f \notin \sigma$.

In such a case the i th fault cannot have occurred, because none of the firing sequences consistent with the observation contains fault transitions in T_f^i .

- $\Delta(w, T_f^i) = 1$ if:

- (i) there exist $\sigma \in \mathcal{S}(w)$ and $t_f \in T_f^i$ such that $t_f \in \sigma$ but

- (ii) for all $(\sigma_o, \sigma_u) \in \hat{\mathcal{J}}(w)$ and for all $t_f \in T_f^i$ it holds that $t_f \notin \sigma_u$.

In such a case a fault transition of the i th class may have occurred but is not contained in any justification of w .

- $\Delta(w, T_f^i) = 2$ if there exist $(\sigma_o, \sigma_u), (\sigma'_o, \sigma'_u) \in \hat{\mathcal{J}}(w)$ such that

- (i) there exists $t_f \in T_f^i$ such that $t_f \in \sigma_u$;

- (ii) for all $t_f \in T_f^i, t_f \notin \sigma'_u$.

In such a case a fault transition in the i th class is contained in one (but not in all) justification of w .

- $\Delta(w, T_f^i) = 3$ if for all $\sigma \in \mathcal{S}(w)$ there exists $t_f \in T_f^i$ such that $t_f \in \sigma$.

In such a case the i th fault must have occurred, because all firable sequences consistent with the observation contain at least one fault transition in the i th class. ■

A systematic procedure has been given in [2], [14] to compute the above diagnosis states that is not recalled here for the sake of brevity.

Example 4.3: Let us consider again the PN in Fig. 2, where $T_f = \{\varepsilon_5, \varepsilon_7\}$.

Let $w = ab$. In such a case it is $\Delta(w, T_f) = 2$. In fact, the j-vector $y' = [1 \ 0 \ 0 \ 0 \ 0]^T$ does not contain fault transitions, while $y'' = [1 \ 0 \ 1 \ 1 \ 1]^T$ contains $\varepsilon_7 \in T_f$. ■

V. DECENTRALIZED DIAGNOSIS

In this section we present the main contributions of the paper. In particular, we introduce two different protocols to solve the decentralized diagnosis problem introduced in Section III.

A. Diagnosis under Protocol 1

Protocol 1 is based on the following very simple rules.

Let σ be the sequence that has occurred and $w_j = \mathcal{L}_j(\sigma)$ be the observation of site $j \in \mathcal{J}$. We denote as $\Delta_{j,i} = \Delta(w_j, T_f^i)$ the diagnosis state of site j wrt T_f^i .

- 1) The diagnosis state $\bar{\Delta}_i$ of the coordinator relative to each T_f^i is initially undefined.
- 2) If there exists a site j such that $\Delta_{j,i} = 3$ for some $i \in \mathcal{F}$, then the site j communicates this information to the coordinator; otherwise it remains silent.

3) When the coordinator receives some information relative to a fault class i , then it sets $\bar{\Delta}_i = 3$. This means that a fault in T_f^i has been detected.

A decentralized diagnoser following Protocol 1 satisfies the following important property. Note that in the following we denote as Δ_i^* the diagnosis state relative to the i -th fault class computed using the centralized approach with set of observable transitions T_o summarized in the previous section, that is assumed as a target.

Proposition 5.1: The coordinator based on Protocol 1 never produces false alarms, namely if $\bar{\Delta}_i = 3$, then $\Delta_i^* = 3$ as well.

Proof: If the coordinator diagnosis state is $\bar{\Delta}_i = 3$, it means that there exists at least one site $j \in \mathcal{J}$ such that $\Delta_{j,i} = 3$. Now, by eq. (1) it is $T_{u,j} \supseteq T_u$. As a consequence, all the justifications that are admissible for the centralized diagnoser are also admissible for the j -th site. However, there may exist other justifications that are admissible for the j -th site while they are not admissible for the centralized diagnoser. This implies that if $\Delta_{j,i} = 3$ then all the justifications computed by the j -th site contain fault transitions in T_f^i , then for sure any subset of such justifications (including the set of justifications computed by the centralized diagnoser) contains fault transitions in T_f^i , thus proving the statement. \square

It is important to note that it may happen that the centralized diagnosis state is $\Delta_i^* = 3$, while the coordinator under Protocol 1 is silent because the diagnosis state of all the sites are equal to 2 wrt fault class T_f^i .

Example 5.2: Let us consider the Petri net system in Fig. 3 containing only one fault transition t_f . Assume that the diagnosis is performed according to Protocol 1 by two sites whose sets of observable labels (alphabets) are equal to $L_1 = \{a, c\}$ and $L_2 = \{b, c\}$, respectively.

Assume that the sequence $t_f t_3 t_4 t_5^k$ fires, where k is an arbitrary integer number.

A centralized diagnoser whose alphabet is $L = \{a, b, c\}$ observes the word $w = bac^k$ that has only the justification $\sigma_u = t_f$. Thus its diagnosis state is set equal to 3.

The word observed by site 1 is $w_1 = ac^k$ to which correspond two different justifications $\sigma'_{u,1} = t_f t_3$ and $\sigma''_{u,1} = t_2$, one containing the fault and the other one not. Thus its diagnosis state is set equal to 2.

Similarly, the word observed by site 2 is $w_2 = bc^k$ to which correspond two different justifications, one containing the fault and the other one not, namely, $\sigma'_{u,2} = t_f t_4$ and $\sigma''_{u,2} = t_1$. Thus its diagnosis state is set equal to 2.




Fig. 3. Petri net system considered in Example 5.2.

According to Protocol 1 the two sites remain silent so the coordinator does not detect the fault. ■

Let us now discuss diagnosability. The following result obviously holds.

Corollary 5.3: If a system is diagnosable in the decentralized framework, then it is also diagnosable in the centralized framework. ■

Clearly, the other sense of the implication does not hold. However, in the case of diagnosis performed using Protocol 1 the following result can be proved.

Proposition 5.4: The system is diagnosable wrt the decentralized approach based on Protocol 1 iff for any fault class $i \in \mathcal{F}$ there exists at least one site $j \in \mathcal{J}$ such that the system is diagnosable wrt the centralized approach with set of observable transitions $T_{o,j}$ and wrt that fault class.

Proof: For simplicity, with no loss of generality we assume that there is only one fault class. Let us prove separately the if and only if statements.

(If) If there exists one site $j \in \mathcal{J}$ such that the system is diagnosable wrt the centralized approach with set of observable transitions $T_{o,j}$, due to Assumptions **A1** and **A2**, this means that the j -th site reconstructs for sure the occurrence of a fault in a finite number of steps. Therefore its diagnosis state becomes equal to 3 after a finite number of transitions firings, as well as the diagnosis state of the coordinator.

(Only if) We prove this by contradiction. Assume that the system is diagnosable wrt the centralized approach with set of observable transitions $T_{o,j}$, but not wrt the decentralized approach. This means that even if a fault is contained in all the justifications computed assuming $T_{o,j}$ as the set of observable transitions, then $\Delta^* = 3$ while $\bar{\Delta}_j \neq 3$. But this leads to a contradiction because, by Assumption **A1**, being the set of transitions observable to the centralized diagnoser

equal to $T_{o,j}$, the set of justifications is the same in the two cases. \square

B. Diagnosis under Protocol 2

Protocol 2 is a generalization of Protocol 1. It is still based on the idea that a site communicates its diagnosis state if and only if it is equal to 3, otherwise it remains silent. However, in this case it also transmits its set of j-vectors. On the basis of this information, the coordinator polls a certain number of sites and makes a refinement of the set of j-vectors. Such a refinement is then used by the local site to recompute its diagnosis states. This in general leads to an improvement of the performance of the decentralized diagnoser.

To define in a clear and concise way such a protocol, let us introduce some preliminary definitions.

- Let $\mathcal{J}_l = \{k \in \mathcal{J} \mid l \in L_k\}$ be the set of sites that are capable of observing label l .
- Given a site j and an observed word w_j ,

$$\mathcal{I}(j, w_j) = \{l \in L \mid \exists y \in Y_{min}(M_0, w_j), \\ y(t) > 0 \wedge \mathcal{L}(t) = l\}$$

is the set of labels relative to transitions that appear in at least a j-vector of the j -th module.

- Let $|w_k|_l$ be the number of occurrences of label l in the observation w_k .
- Given an observation w_k from site k , a label l , and a j-vector y ,

$$\beta_k(l, y) = |w_k|_l - \sum_{t:\mathcal{L}(t)=l} y(t)$$

is the difference between the number of times the site k has observed l and the number of times a transition labeled l appears in y .

Based on the above definitions, the main steps of the decentralized procedure based on Protocol 2 can be summarized as follows.

- 1) The diagnosis state $\bar{\Delta}_i$ of the coordinator relative to each T_f^i is initially undefined.
- 2) If $\Delta_{j,i} = \Delta(w_j, T_f^i) = 3$ for some $j \in \mathcal{J}$ and some $i \in \mathcal{F}$, then the j -th site transmits to the coordinator its diagnosis state together with its set of j-vectors.
- 3) For any label $l \in \mathcal{I}(j, w_j)$ the coordinator polls one site $k \in \mathcal{J}_l \setminus \{j\}$.
- 4) The k -th site transmits to the coordinator the value of $|w_k|_l$.

- 5) If $\beta_k(l, y) < 0$ for a vector $y \in Y_{min}(M_0, w_j)$, then the coordinator removes the vector y from the set of j -vectors $Y_{min}(M_0, w_j)$ relative to the j -th site.
- 6) As a result of this process of refinement, the coordinator computes a new set $Y'_{min}(M_0, w_j)$ that is communicated to the j -th site.
- 7) The j -th site recomputes its diagnosis states according to the new set $Y'_{min}(M_0, w_j)$ and if some of them are equal to 3, communicates it to the coordinator, otherwise it keeps silent.

The refinement of $Y_{min}(M_0, w_j)$ is based on the following very simple fact. If $Y_{min}(M_0, w_j)$ contains a j -vector that assumes a certain number of occurrences of l , but this number is not consistent with the observation of a site that is capable of observing l , then for sure such a justification is unfeasible. Therefore, if $\beta_k(l, y) < 0$ for a certain label l and a certain j -vector $y \in Y_{min}(M_0, w_j)$, then y should be removed from $Y_{min}(M_0, w_j)$. In fact, this means that the justification relative to j -vector y assumes a number of occurrences of l that is greater than the real number, that is perfectly known by the k -th site. On the contrary, if $\beta_k(l, y) \geq 0$ it means that the j -vector y is feasible. In particular, if $\beta_k(l, y) = 0$ it means that the justification contains all the occurrences of label l . The case of $\beta_k(l, y) > 0$ is relative to a feasible situation as well. It means that the justification relative to y does not contain all the occurrences of l ; thus the rest of transitions labeled l , up to the value $|w_k|_l$, have fired after the justification and the observation w_j .

The refinement process has in general positive effects on diagnosis as shown by the following example.

Example 5.5: Let us consider the Petri net system in Fig. 4. Assume that there are two fault classes: $T_f^1 = \{t_{f,1}^I, t_{f,1}^{II}\}$, $T_f^2 = \{t_{f,2}\}$.

Assume that the net is locally diagnosed by two sites whose sets of observable transitions are $T_{o,1} = \{t_3, t_6\}$ and $T_{o,2} = \{t_1, t_2, t_4, t_5, t_6\}$, respectively. This implies that $L_1 = \{a, c\}$, $L_2 = \{b, c\}$, $\mathcal{J}_a = \{1\}$, $\mathcal{J}_b = \{2\}$ and $\mathcal{J}_c = \{1, 2\}$.

Assume that the sequence $\sigma = t_{f,1}^I t_1 t_{f,2} t_2$ fires, thus $w = \mathcal{L}(\sigma) = bb$. The first site observes the empty string ε , i.e., $w_1 = \varepsilon$, while the second site observes the word $w_2 = bb$.

Due to these observations, the diagnosis states of the first site are $\Delta_{1,1} = 1$ and $\Delta_{1,2} = 1$, relative to the first and the second fault class respectively. In fact, transitions from both fault classes may have fired at the initial marking without the firing of any transition labeled either a or c .




Fig. 4. Petri net system considered in Example 5.5.

The diagnosis states of the second site are $\Delta_{2,1} = 3$ and $\Delta_{2,2} = 2$, respectively. In fact, the set of justifications of w_2 includes the following sequences: $\sigma'_{u,2} = t_{f,1}^I t_{f,2}$, $\sigma''_{u,2} = t_3 t_{f,1}^{II}$, i.e., both the justifications contain a transition in T_f^1 , while only one of them contains a transition in T_f^2 .

Therefore, the second site communicates $\Delta_{2,1} = 3$ to the coordinator who sets its diagnosis state relative to T_f^1 to $\bar{\Delta}_1 = 3$. The firing of one transition in T_f^1 is thus detected both using Protocol 1 and Protocol 2.

However, if we use Protocol 1 the firing of $t_{f,2}$ is not detected because both sites are silent wrt the second fault class. On the contrary, if we use Protocol 2 the firing of $t_{f,2}$ is detected.

In fact, according to Protocol 2, site 2 also communicates its set of j-vectors to the coordinator that is equal to $Y_{min}(M_0, w_2) = \{y'_2, y''_2\}$, where y'_2 is the firing vector relative to $\sigma'_{u,2} = t_{f,1}^I t_{f,2}$, while y''_2 is the firing vector relative to $\sigma''_{u,2} = t_3 t_{f,1}^{II}$.

Since $\mathcal{I}(2, w_2) = \{a\}$ and $\mathcal{J}_a = \{1\}$, the coordinator polls site 1 to know the number of symbols a it has observed. Since $|w_1|_a = 0$, then $\beta_1(a, y''_2) = 0 - 1 < 0$. It means that j-vector y''_2 can be confuted and removed from $Y_{min}(M_0, w_2)$. The refined set of j-vectors is $Y'_{min}(M_0, w_2) = \{y'_2\}$ thus $\Delta_{2,2}$ is updated to 3 and consequently $\bar{\Delta}_2 = 3$ allowing also the detection of $t_{f,2}$. ■

Remark 5.6: Since events occur in an asynchronous way, it can obviously happen that the value of $|w_k|_l$ transmitted by the polled sites to the coordinator is affected by some delay. As a result of this the coordinator receives a value $|w_k|'_l > |w_k|_l$ because during such a delay other transitions labeled l may have fired. This implies that the value of $\beta_k(l, y)$ may be greater than the correct one. In particular, it may occur that a negative value of $\beta_k(l, y)$ becomes null or even positive, thus certain j-vectors that should be rejected, are considered as feasible. However such a delay may never cause a feasible j-vector to be rejected. □

The following propositions can be stated.

Proposition 5.7: The coordinator based on Protocol 2 never produces false alarms, namely if $\bar{\Delta}_i = 3$, then $\Delta_i^* = 3$ as well.

Proof: A formal proof can be obtained using the same arguments of Proposition 5.1. Thus it is omitted for the sake of brevity. \square

Proposition 5.8: All sets of j-vectors obtained as the result of a refinement carried out according to the rules of Protocol 2, are not empty, i.e., $Y'_{min}(M_0, w_j) \neq \emptyset$ for all $j \in \mathcal{J}$ that perform a refinement of $Y_{min}(M_0, w_j)$.

Proof: Follows from the fact that the set $Y_{min}(M_0, w_j)$ contains certainly the j-vector \bar{y} that corresponds to the word that has actually fired, plus eventually other vectors. Using the rules of Protocol 2, some of these j-vectors will be confuted, but certainly it will not be \bar{y} , therefore $\bar{y} \in Y'_{min}(M_0, w_j)$, thus proving the statement. \square

Proposition 5.9: The system is diagnosable wrt the decentralized approach based on Protocol 2 if for any fault class $i \in \mathcal{F}$ there exists at least one site $j \in \mathcal{J}$ such that the system is diagnosable wrt the centralized approach with set of observable transitions $T_{o,j}$ and wrt that fault class.

Proof: This result can be proved using the same arguments in the proof of the *if* statement of Proposition 5.4 \square

The above proposition only provides a sufficient condition for diagnosability. In fact let us consider for the sake of simplicity only one fault class. It may happen that the system is not diagnosable in a centralized framework wrt all $T_{o,j}$ ($j \in \mathcal{J}$), while it is diagnosable in a decentralized framework using ν sites whose sets of observable transitions are equal to $T_{o,j}$ ($j \in \mathcal{J}$).

This is the case of the Petri net system considered in Example 5.5. In fact, both the centralized diagnosers observing $T_{o,1} = \{t_3, t_6\}$ and $T_{o,2} = \{t_1, t_2, t_4, t_5, t_6\}$ are not able to detect the occurrence of $t_{f,2}$ if the sequence $\sigma = t_{f,1}^l t_1 t_{f,2} t_2 t_6^k$ fires, where k is an arbitrary integer number. On the contrary, as shown in Example 5.5, the decentralized diagnoser based on Protocol 2 detects the occurrence of $t_{f,2}$ after a sequence that is a prefix of σ .

We also observe that, as in the case of Protocol 1, it may happen that the centralized diagnosis state is $\Delta_i^* = 3$ while the coordinator under Protocol 2 is silent. The following example clarifies this.




Fig. 5. Petri net system considered in Example 5.10.

Example 5.10: Let us consider the net system in Fig. 5, having a single fault transition t_f . The net is locally diagnosed by two sites whose alphabets are equal to $L_1 = \{a, c\}$ and $L_2 = \{b, c\}$, respectively.

Assume that the sequence $\sigma = t_f t_1 t_4$ fires, thus $w_1 = a$ and $w_2 = b$.

The set of j-vectors relative to the first site is $Y_{min}(M_0, w_1) = \{y'_1, y''_1\}$ where y'_1 is the firing vector relative to the justification $\sigma'_{u,1} = t_f$, while y''_1 is the firing vector relative to $\sigma''_{u,1} = \varepsilon$. The set of j-vectors relative to the second site is $Y_{min}(M_0, w_2) = \{y'_2, y''_2\}$ where y'_2 and y''_2 are relative respectively to justifications $\sigma'_{u,2} = t_f t_1$ and $\sigma''_{u,2} = t_2 t_3$. Hence both sites have a diagnosis state equal to 2.

On the contrary, in a centralized framework, being $L = \{a, b, c\}$ and consequently $w = ab$, the diagnosis state is equal to 3 and the firing of t_f is detected. In fact the only justification of w is $\sigma_u = t_f$. ■

We conclude the paper with the following remark.

Remark 5.11: Assume, for simplicity of explanation, that there is only one fault class.

According to the proposed protocols the coordinator may either be in a fault state or it may be silent. If the coordinator is silent, the fault may either have occurred or not.

If we also want to characterize the situation in which the occurrence of a fault can be excluded for sure, both protocols can be modified as follows. Three different states are defined for the coordinator, e.g., F (fault), U (uncertain) and N (no fault). The sites communicate their diagnosis state to the coordinator even if it is equal to 0. If the coordinator receives one 0, then it sets to N its fault state; if it receives one 3 then it sets to F its state; otherwise its state is equal to U . ■

VI. CONCLUSIONS AND FUTURE WORKS

In this paper we addressed the problem of designing a decentralized diagnoser for Petri nets. We assume that the system is monitored by a set of local sites: each site knows the structure of the net and the initial marking of the system but observes its evolution with a different mask. Diagnosis is performed locally using a diagnosis approach we previously introduced in the centralized framework. Two different protocols are proposed to determine how a central coordinator elaborates the global diagnosis states. The problem of diagnosability is also addressed and the advantages/disadvantages of the two protocols are discussed.

Our future work will be that of investigating if the performance of the decentralized diagnoser, and its diagnosability properties can be improved, if the sites communicate with the coordinator also in the case of diagnosis state equal to 2, or 1. The problem of determining a technique to test diagnosability in the case of decentralized systems will also be addressed.

Finally, while in this paper we assumed that the sites and their observation masks are given, we will also consider the case in which their definition can be seen as the result of an optimization problem, whose main goal is that of obtaining performances in terms of diagnosis (and diagnosability) that are close as possible to those of the centralized diagnoser.

REFERENCES

- [1] F. Basile, P. Chiacchio, and G. D. Tommasi, "An efficient approach for online diagnosis of discrete event systems," *IEEE Trans. Automatic Control*, vol. 54, no. 4, pp. 748–759, 2008.
- [2] M. Cabasino, A. Giua, and C. Seatzu, "Diagnosis of discrete event systems using labeled Petri nets," in *Proc. 2nd IFAC Workshop on Dependable Control of Discrete Systems (Bari, Italy)*, Jun. 2009.
- [3] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, "Diagnosability of discrete-event systems," *IEEE Trans. Automatic Control*, vol. 40 (9), pp. 1555–1575, 1995.
- [4] Y. Wu and C. Hadjicostis, "Algebraic approaches for fault identification in discrete-event systems," *IEEE Trans. Robotics and Automation*, vol. 50, no. 12, pp. 2048–2053, 2005.
- [5] S. H. Zad, R. Kwong, and W. Wonham, "Fault diagnosis in discrete-event systems: framework and model reduction," *IEEE Trans. Automatic Control*, vol. 48, no. 7, pp. 1199–1212, Jul. 2003.
- [6] R. Boel and J. van Schuppen, "Decentralized failure diagnosis for discrete-event systems with costly communication between diagnosers," in *Proc. WODES'02: 6th Work. on Discrete Event Systems (Zaragoza, Spain)*, Oct. 2002, pp. 175–181.
- [7] O. Contant, S. Lafortune, and D. Teneketzis, "Diagnosability of discrete event systems with modular structure," *Discrete Event Dynamic Systems*, vol. 16, no. 1, pp. 9–37, 2006.
- [8] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated decentralized protocols for failure diagnosis of discrete-event systems," *Discrete Events Dynamic Systems*, vol. 10, no. 1, pp. 33–86, 2000.

- [9] R. Su, W. Wonham, J. Kurien, and X. Koutsoukos, "Distributed diagnosis for qualitative systems," in *6th International Workshop on Discrete Event Systems, Zaragoza*, 2002, pp. 169–174.
- [10] Y. Wang, T.-S. Yoo, and S. Lafortune, "Diagnosis of discrete event systems using decentralized architectures," *Discrete Event Dynamic Systems*, vol. 17, no. 2, 2007.
- [11] A. Benveniste, E. Fabre, S. Haar, and C. Jard, "Diagnosis of asynchronous discrete event systems, a net unfolding approach," *IEEE Trans. Automatic Control*, vol. 48, no. 5, pp. 714–727, May 2003.
- [12] S. Genc and S. Lafortune, "Distributed diagnosis of place-bordered Petri nets," *IEEE Trans. on Automation Science and Engineering*, vol. 4, no. 2, pp. 206–219, 2007.
- [13] G. Jiroveanu and R. K. Boel, "A distributed approach for fault detection and diagnosis based on time Petri nets," *Mathematics and Computers in Simulation*, vol. 70, no. 5, 2006.
- [14] M. Cabasino, A. Giua, and C. Seatzu, "Fault detection for discrete event systems using Petri nets with unobservable transitions," *Automatica*, (Preliminary accepted).
- [15] C. Cassandras and S. Lafortune, *Introduction to discrete event systems, Second Edition*. Springer, 2007.