A new protocol for the decentralized diagnosis of labeled 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, first we present a protocol that defines the communication rules between the sites and the coordinator. Secondly, we prove that this protocol does not produce false alarms. Moreover, we give sufficient conditions for diagnosability based on the notion of failure ambiguous strings. Finally, we compare the protocol here presented with two other protocols that we presented in a previous work.

Published as:

M.P. Cabasino, A. Giua, A. Paoli, C. Seatzu, "A new protocol for the decentralized diagnosis of labeled Petri nets," *WODES10: 10th Int. Work. on Discrete Event Systems* (Berlin, Germany), Aug-Sep 2010.

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

M.P. Cabasino, A. Giua and C. Seatzu are with the Dept. of Electrical and Electronic Engineering, University of Cagliari, Italy, e-mail: {cabasino, giua, seatzu}@diee.unica.it.

A. Paoli is with the Department of Electronic, Computer Science and Systems, University of Bologna, Italy, e-mail: 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 contributes have been presented for discrete event systems in the centralized framework (1; 2; 3; 4; 5). Due to the intrinsic distributed nature of real systems, distributed diagnosis techniques, that take advantage of the natural decompositions of a modular system, have been proposed both in the automata framework (6; 7; 8; 9; 10) and in the Petri net (PN) framework (11; 12; 13; 14).

In particular, (11) solves a problem of alarm supervision in telecommunication networks. They use an unfolding approach and restrict their attention to safe PNs. (13) proposes a diagnoser on the basis of a modular approach that performs the diagnosis of faults in each module. Subsequently, the diagnosers recover the monolithic diagnosis information obtained when all the modules are combined into a single module that preserves the behavior of the underlying modular system. A communication system connects the different modules and updates the diagnosis information. In (12) is proposed an algorithm for the model based design of a distributed protocol for fault detection and diagnosis for very large systems. The overall process is modeled as different timed PN models that interact with each other via guarded transitions that become enabled only when certain conditions are satisfied. Different local agents receive local observation as well as messages from neighboring agents. Each agent estimates the state of the part of the overall process for which it has model and from which it observes events by reconciling observations with model based predictions.

In (14) we presented two different protocols for decentralized diagnosis of labeled Petri nets based on a particular architecture, that is the same we consider in this paper. In particular, we assume that the system can be observed by different local sites that have the perfect knowledge of the net system, but observe its evolution with different masks. On the basis of its own observations, each site performs diagnosis locally.

Here we present a third protocol that defines the communication rules between the local sites and the coordinator. It differs from the ones defined in (14) because it leads to more accurate diagnosis. The price to pay for this improvement in the performances is that a larger amount

of information should be exchanged between the sites and the coordinator. We prove that this protocol, as well as those introduced in (14), never produces false alarm.

Furthermore, we analysis diagnosability. To this aim, we recall the definition of failure ambiguous strings, and show that the absence of such kind of strings is only a sufficient condition for the diagnosability of a Petri net system using Protocol 3.

We conclude this section observing that both the problem formulation and the objectives considered in (11) are significantly different from those in this paper. More strict analogies exist between our approach and the approaches of (13) and (12). However, also in this case there exist a main difference that can be summarized as follows. In these works the authors assume the PN divided into different sub-modules or sites: each site is modeled by a different subset of places and transitions and can interact with the other sites via a restricted interface consisting in bordered places (13) or guard transitions (12). On the contrary, in our approach each site has the perfect knowledge of the whole PN system but observes the system with a different observation mask and no special interfaces are required.

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, $Pre : P \times T \to \mathbb{N}$ and $Post : P \times T \to \mathbb{N}$ 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 $M : P \to \mathbb{N}$ 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'. Given a sequence $\sigma \in T^*$ we write $t \in \sigma$ to denote that a transition t is contained in σ.

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 $\pi : T^* \to \mathbb{N}^n$ 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 no 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* $\mathcal{L} : T \to L \cup \{\varepsilon\}$ assigns to each transition a symbol from a given alphabet L or the empty string ε. We denote as $\mathcal{L}^{-1}$ the inverse operator of $\mathcal{L}$. 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). We define the *projection over T_o (projection over T_u)* $P_o : T^* \to T_o^*$ ($P_u : T^* \to T_u^*$) as: (i) $P_o(\varepsilon) = \varepsilon$ ($P_u(\varepsilon) = \varepsilon$); (ii) for all $\sigma \in T^*$ and $t \in T$, $P_o(\sigma t) = P_o(\sigma)t$ if $t \in T_o$ ($P_u(\sigma t) = P_u(\sigma)t$ if $t \in T_u$), and $P_o(\sigma t) = P_o(\sigma)$ ($P_u(\sigma t) = P_u(\sigma)$) otherwise.

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

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

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 $i \in \mathcal{F} = \{1, \ldots, r\}$, that model different fault classes. The transition set $T_{reg} = T_u \setminus T_f$ represents the set of unobservable, but regular, transitions.

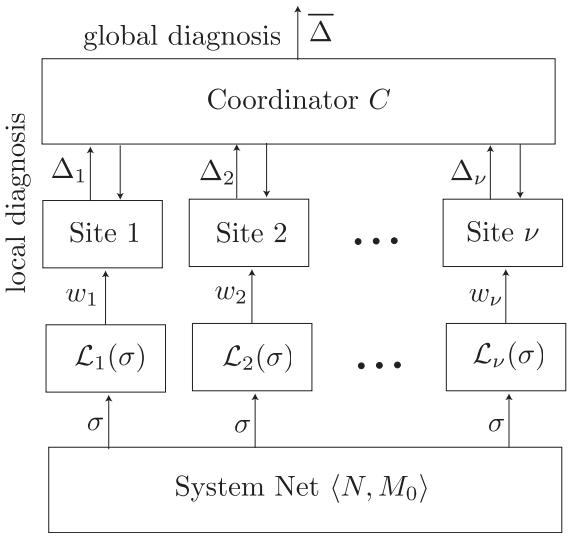

Fig. 1. The decentralized diagnosis architecture.

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 $\mathcal{J} = \{1, \ldots, \nu\}$ 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

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

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

$$\mathcal{L}_| : T \to L_j \cup \{\varepsilon\} \tag{2}$$

the labeling function associated to the j-th site and as

$$\bar{\mathcal{L}} : T \rightarrow L \cup \{\varepsilon\} \tag{3}$$

the labeling function associated to the centralized system. Finally, $w_j = \mathcal{L}_j(\sigma)$ denoted 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 protocol[1]. 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.

A5 The system does not enter a deadlock after the firing of any fault transition.

In this paper we also investigate the issue of *diagnosability*.

***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. ■

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

***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:

- $\bar{\mathcal{L}}(\sigma_\infty) = \bar{\mathcal{L}}(\sigma_\in)$,
- $\sigma_1 \in (T \setminus T_f^i)^*$,
- $\exists$ 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 (5) 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 $\bar{\mathcal{L}}(\sigma_\imath) = \sqsupseteq$. We call *justification of w* a 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 $Y_{min}(M_0, w)$ the set of firing vectors relative to justifications of w.
  The generic element $y \in Y_{min}(M_0, w)$ is called *j-vector*.
- Finally, we denote as

$$\begin{aligned}\hat{\mathcal{J}}(w) \;\; = \{\; &(\sigma_o, \sigma_u),\; \sigma_o \in T_o^*,\; \bar{\mathcal{L}}(\sigma_\imath) = \sqsupseteq,\\ &\sigma_u \in T_u^* \;\;|\\ &[\exists \sigma \in \mathcal{S}(w) \;:\; \sigma_o = P_o(\sigma),\; \sigma_u = P_u(\sigma)] \wedge\\ &[\nexists \sigma' \in \mathcal{S}(w) :\; \sigma_o = P_o(\sigma'),\; \sigma'_u = P_u(\sigma') \wedge\\ &\qquad \pi(\sigma'_u) \lneq \pi(\sigma_u)]\}\end{aligned}$$

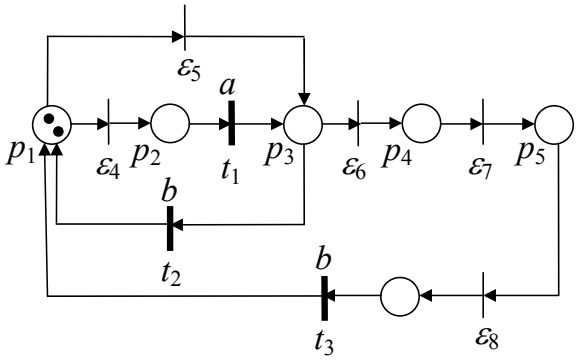

Fig. 2. The Petri net system considered in Examples 4.1 and 4.3.

the set of couples (sequence $\sigma_o \in T_o^*$ with $\bar{\mathcal{L}}(\sigma_o) = \sqsupseteq$ - 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 $\bar{\mathcal{L}}(\sqcup_\infty) = \dashv$ and $\bar{\mathcal{L}}(\sqcup_\in) = \bar{\mathcal{L}}(\sqcup_\ni) = \lfloor$.

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 $\Delta : L^* \times \{T_f^1, T_f^2, \ldots, T_f^r\} \to \{0, 1, 2, 3\}$ that associates to each observation w and to each fault class T_f^i, $i \in \mathcal{F}$, 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 ith 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 ith 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 ith 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 ith fault must have occurred, because all firable sequences consistent with the observation contain at least one fault transition in the ith class. ■

A systematic procedure has been given in (5) 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 using Protocol 3

Protocol 3 is based on the idea that a site communicates its diagnosis state if and only if it is equal either to 3 or to 2, otherwise it remains silent. Each site transmits not only the diagnosis state but also 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 sites to recompute their diagnosis states for all fault classes. 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 a set of j-vectors $Y_j = Y_{min}(M_0, w_j)$,

$$\mathcal{I}(j, Y_j) = \{l \in L \ \mid \ \exists\ y \in Y_j \ \wedge \exists\ t \in T \setminus T_{o,j} \ : \\ y(t) > 0 \wedge \mathcal{L}(t) = l\}$$

  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,

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

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 3 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 $\Delta_{j,i} = \Delta(w_j, T_f^i) = \{2, 3\}$ 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 $l \in \mathcal{I}(j, Y_j)$ the coordinator polls any site $k \in \mathcal{J}_l \setminus \{j\}$ (if $\mathcal{J}_l \setminus \{j\}$ is not empty).
4) The k-th site transmits to the coordinator the value of $|w_k|_l$.
5) If $\beta_k(w_k, l, y) < 0$ for a vector $y \in Y_j$, then the coordinator removes the vector y from the set of j-vectors Y_j relative to the j-th site.
6) As a result of this process of refinement, the coordinator computes a new set Y_j' that is communicated to the j-th site.
7) The j-th site recomputes its diagnosis states according to the new set Y_j' and if some of them are equal to 3, communicates it to the coordinator, otherwise it keeps silent.

The refinement of Y_j is based on the following very simple fact. If Y_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 $\beta_k(w_k, l, y) < 0$ for a certain label l and a certain j-vector $y \in Y_j$, then y should be removed from Y_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 $\beta_k(w_k, l, y) \geq 0$ it means that the j-vector y is compatible with the observation of the k-th site. In particular, if $\beta_k(w_k, l, y) = 0$ it means that the justification contains all the occurrences of label l. The case of $\beta_k(w_k, l, y) > 0$ is relative to a possible 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.

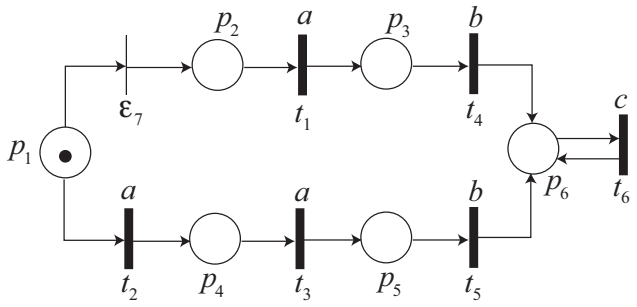

Fig. 3. Petri net system considered in Example 5.1.

***Example** 5.1:* Let us consider the Petri net in Fig. 3 where $T_u = T_f = \{\varepsilon_7\}$. The net is locally diagnosed by two sites whose set of observable transitions is $T_{o,1} = \{t_1, t_3, t_6\}$ and $T_{o,2} = \{t_4, t_5, t_6\}$, respectively. This implies that $L_1 = \{a, c\}$, $L_2 = \{b, c\}$, $\mathcal{J}_a = \{1\}$, $\mathcal{J}_b = \{2\}$ and $\mathcal{J}_c = \{1, 2\}$.

Let us assume that the sequence $\sigma = \varepsilon_7 t_1 t_4$ fires, thus $w_1 = a$ and $w_2 = b$.

The set of j-vectors for the first site is $Y_{min}(M_0, w_1) = Y_1 = \{y_1', y_1''\}$, where $y_1' = \vec{0}$ and $y_1'' = \pi(\varepsilon_7)$, while for the second site is $Y_{min}(M_0, w_2) = Y_2 = \{y_2', y_2''\}$, where $y_2' = \pi(\varepsilon_7 t_1)$ and $y_2'' = \pi(t_2 t_3)$. Hence both sites have a diagnosis state equal to 2.

Both the sites communicate their diagnosis state and their set of j-vectors to the coordinator. Now, $\mathcal{I}(1, Y_1) = \emptyset$ but $\mathcal{I}(2, Y_2) = \{a\}$ and $\mathcal{J}_a = \{1\}$. Thus the coordinator polls site 1 to know the number of label a it has observed. Since $|w_1|_a = 1$, then $\beta_1(w_1, a, y_2') = 1 - 1 = 0$ and $\beta_1(w_1, a, y_2'') = 1 - 2 < 0$. This means that the j-vector $y_2'' = \pi(t_2 t_3)$ can be confuted and then removed from Y_2. The redefined set of j-vectors for site 2 is $Y'_{min}(M_0, w_2) = \{y_2'\}$ and it is communicated by the coordinator to the site 2. Site 2 recomputes its diagnosis state that is now equal to 3. Thus $\Delta_2 = 3$ is communicated to the coordinator and consequently $\bar{\Delta} = 3$ and the fault ε_7 is detected.

■

Let us finally prove the following important property of Protocol 3.

***Proposition** 5.2:* The coordinator under Protocol 3 does not produce any false alarm, 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 $\Delta_{j,i} = 3$. It may happen than either $\Delta_{j,i} = 3$ as soon as the diagnosis state is computed or that $\Delta_{j,i}$ becomes equal to 3 after the confutation procedure.

Let us analyze these two situations separately. Now, for the first case, 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.

For the second case, the reduction of the cardinality of the sets of j-vectors relative to certain sites cannot produce false alarm as well. In fact, by definition such a reduction consists in only removing those j-vectors that for sure are not feasible, because they are not consistent with the observations of other sites. Thus in both situations false alarms cannot be produced. □

VI. DIAGNOSABILITY ANALYSIS

The first important step when analyzing the decentralized diagnosability of a PN system is that of detecting the presence of particular strings, called *failure ambiguous strings*. This notion has been firstly introduced in (6) in the framework of automata. In particular, in (6) the authors assume that the decentralized diagnoser only includes two sites. In (16) we generalized such a definition to PNs and consider the general case of an arbitrary number ν of sites.

***Definition** 6.1:* Consider a net system $\langle N, M_0 \rangle$ monitored by a set $\mathcal{J} = \{1, \ldots, \nu\}$ of sites. Let $T_{o,j} \subseteq T_o$ be the set of locally observable transitions for the generic site $j \in \mathcal{J}$. Finally, let $T_f^i \subseteq T_f$ be the generic i-th fault class, with $i \in \mathcal{F}$.

A string $\sigma \in T^*$ of arbitrary length, such that $t_f \in \sigma$ for at least one $t_f \in T_f^i$, is said to be *failure ambiguous* wrt the above set of sites and wrt the fault class T_f^i, if the following two conditions are verified:

(a) $\mathcal{L}_j^{-1}(\mathcal{L}_j(\sigma)) \cap (T \setminus T_f^i)^* \neq \emptyset \quad \forall j \in \mathcal{J}$;

(b) $\bar{\mathcal{L}}^{-\infty}(\bar{\mathcal{L}}(\sigma)) \cap (\mathcal{T} \setminus \mathcal{T}_{\{}^{\rangle})^* = \emptyset$,

where $\mathcal{L}_j$ and $\bar{\mathcal{L}}$ are defined as in (2), (3), respectively. ■

In simple words, a sequence σ of arbitrary length containing some fault transitions in a fault class i, is failure ambiguous wrt to a set of sites and wrt the i-th fault class, if the word σ is

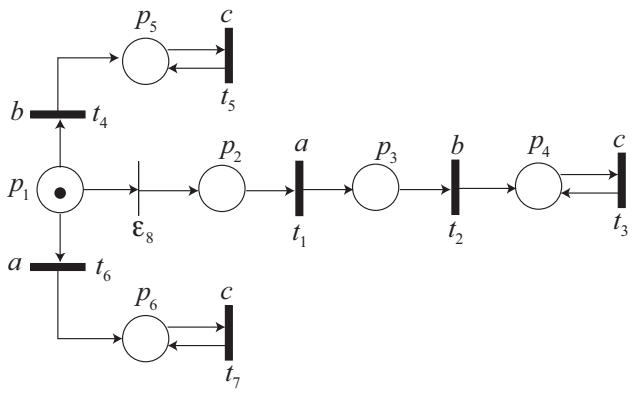

Fig. 4. Petri net system for Example 6.2.

ambiguous for each site $j \in \mathcal{J}$, i.e., it may also be explained by a non faulty word, and the word σ is not ambiguous for the centralized system.

***Example** 6.2:* Let us consider the Petri net system in Fig. 4 which is locally diagnosed by two sites whose alphabets are equal to $L_1 = \{a, c\}$ and $L_2 = \{b, c\}$, respectively. The sequence $\sigma = \varepsilon_8 t_1 t_2 t_3^q$, with $q \in \mathbb{N}$, is failure ambiguous wrt the sites 1 and 2 and wrt to the unique fault class $T_f = \{\varepsilon_8\}$. In fact, $\mathcal{L}_1(\sigma) = \{ac^q\}$ and $\mathcal{L}_1^{-1}(\mathcal{L}_1(\sigma)) = \{\varepsilon_8 t_1 t_2 t_3^q, t_6 t_7^q\}$, thus $\mathcal{L}_1^{-1}(\mathcal{L}_1(\sigma)) \cap (T \setminus T_f)^* = \{t_6 t_7^q\}$; $\mathcal{L}_2(\sigma) = \{bc^q\}$ and $\mathcal{L}_2^{-1}(\mathcal{L}_2(\sigma)) = \{\varepsilon_8 t_1 t_2 t_3^q, t_4 t_5^q\}$ thus $\mathcal{L}_2^{-1}(\mathcal{L}_2(\sigma)) \cap (T \setminus T_f)^* = \{t_4 t_5^q\}$; and $\bar{\mathcal{L}}(\sigma) = \{\dashv \lfloor \rfloor^{\amalg}\}$ and $\bar{\mathcal{L}}^{-\infty}(\bar{\mathcal{L}}(\sigma)) = \{\varepsilon_\forall \sqcup_\infty \sqcup_\in \sqcup_\ni^{\amalg}\}$ thus $\bar{\mathcal{L}}^{-\infty}(\bar{\mathcal{L}}(\sigma)) \cap (\mathcal{T} \setminus \mathcal{T}_{\{}^{\rangle})^* = \emptyset$. ■

In (16) we proved that, if the decentralized architecture is that presented in Section III, regardless of the considered protocol, if a system is diagnosable in a centralized framework with respect to a given fault class, and has no failure ambiguous strings with respect to that class, it is also diagnosable in a decentralized framework. In particular, in (16) we also proposed an efficient method to verify the existence of failure ambiguous strings.

The absence of failure ambiguous strings is only a sufficient condition for the diagnosability in a decentralized framework. Thus, depending on the considered protocol, it may occur that the system is diagnosable in a decentralized framework even in presence of failure ambiguous strings. This is the case of Protocol 3, as illustrated by the following example.

***Example** 6.3:* Let us consider the Petri net system in Fig. 5 where $T_u = T_f = \{\varepsilon_{10}\}$. The net is monitored by two sites whose set of observable transitions is respectively $T_{o,1} = \{t_1, t_3, t_5, t_6, t_7\}$

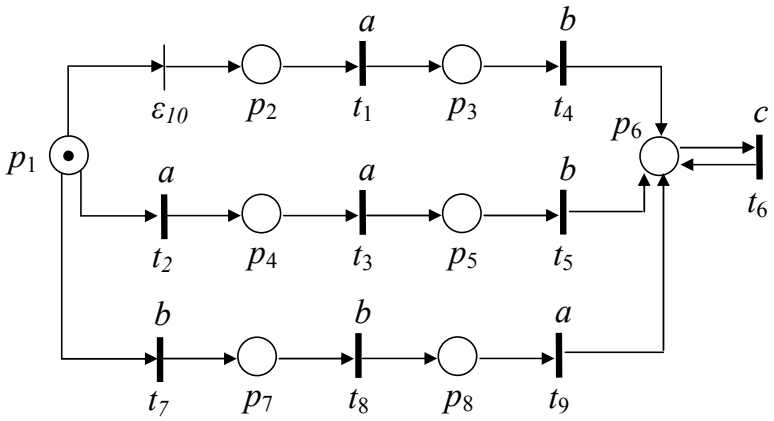

Fig. 5. The Petri net system considered in Example 6.3.

and $T_{o,2} = \{t_2, t_3, t_4, t_5, t_7\}$. This implies that $L_1 = \{a, c\}$, $L_2 = \{b, c\}$, $\mathcal{J}_a = \{1\}, \mathcal{J}_b = \{2\}$ and $\mathcal{J}_c = \{1, 2\}$.

It is easy to verify that all sequences of the form $\sigma = \varepsilon_{10} t_1 t_4 t_6^q$ are failure ambiguous for any $q \in \mathbb{N}$. In fact, $\mathcal{L}_1(\sigma) = \{ac^q\}$ and $\mathcal{L}_1^{-1}(\mathcal{L}_1(\sigma)) = \{\varepsilon_{10} t_1 t_4 t_6^q, t_7 t_8 t_9 t_6^q\}$, thus $\mathcal{L}_1^{-1}(\mathcal{L}_1(\sigma)) \cap (T \setminus T_f)^* = \{t_7 t_8 t_9 t_6^q\}$; $\mathcal{L}_2(\sigma) = \{bc^q\}$ and $\mathcal{L}_2^{-1}(\mathcal{L}_2(\sigma)) = \{\varepsilon_{10} t_1 t_4 t_6^q, t_2 t_3 t_5 t_6^q\}$ thus $\mathcal{L}_2^{-1}(\mathcal{L}_2(\sigma)) \cap (T \setminus T_f)^* = \{t_2 t_3 t_5 t_6^q\}$; and $\bar{\mathcal{L}}(\sigma) = \{\dashv \lfloor \rfloor^{\amalg}\}$ and $\bar{\mathcal{L}}^{-\infty}(\bar{\mathcal{L}}(\sigma)) = \{\varepsilon_\infty \sqcup_\infty \sqcup_\triangle \sqcup_{/}^{\amalg}\}$ thus $\bar{\mathcal{L}}^{-\infty}(\bar{\mathcal{L}}(\sigma)) \cap (\mathcal{T} \setminus \mathcal{T}_{\{}^{\rangle})^* = \emptyset$.

Now, if the two local sites communicate with the coordinator according to Protocol 3, then both of them initially compute a diagnosis state that is equal to 2 after the firing of σ. However, when the confutation procedure is applied, both of them reconstruct the firing of ε_{10}. In particular, the first site observes $w_1 = ac^q$, thus $Y_{\min}(M_0, w_1) = \{\pi(\varepsilon_{10} t_4), \pi(t_7 t_8)\}$ and $\Delta_1 = 2$. Similarly, the second site observes $w_2 = bc^q$ thus $Y_{\min}(M_0, w_2) = \{\pi(\varepsilon_{10} t_1), \pi(t_2 t_3)\}$ and $\Delta_2 = 2$ as well. However, both $\pi(t_7 t_8)$ and $\pi(t_2 t_3)$ are confuted, thus the two diagnosis states become $\Delta_1 = \Delta_2 = 3$ and the fault is diagnosed.

Let us finally observe that, since by inspection it can be verified that the considered family of sequences σ are the only failure ambiguous strings, we can conclude that the system is diagnosable using Protocol 3 even in the presence of failure ambiguous strings. ■

VII. A COMPARISON WITH OUR PREVIOUSLY DEFINED PROTOCOLS

As already mentioned in the Introduction, we presented in (14) two other decentralized protocols, named Protocol 1 and Protocol 2.

Protocol 1 is based on the idea that each local site communicates its diagnosis state to the coordinator if and only if it is equal to 3. No other information is changed, and the coordinator sets its diagnosis state equal to 3 only if it receives a diagnosis state equal to 3 by at least one local site.

Protocol 2 is still based on a confutation procedure, as well as Protocol 3. However, it basically differs from Protocol 3 for the fact that local sites send information to the coordinator if and only if their diagnosis states are equal to 3, while they remain silent if their diagnosis states are 2.

In this section we want to discuss the advantages of using Protocol 3, rather than 1 or 2. Note that obviously Protocol 3 has the disadvantage of requiring a larger amount of information exchanged.

Concerning Protocol 1, the first main issue is that it can be easily proved that using Protocol 1 it can never occur that a system is diagnosable in a decentralized way in the presence of failure ambiguous strings.

On the contrary, it may be the case that a system is diagnosable in a decentralized framework using Protocol 2 even in the presence of failure ambiguous strings if and only if the set of fault transitions is partitioned in at least two fault classes, while it cannot occur in the presence of only one fault class. In fact, if there is a single fault class and there exists at least one failure ambiguous string, for that string the diagnosis states of all sites will be equal to 2, thus under Protocol 2 all sites remain silent and the fault cannot be diagnosed. Note that Protocol 3 does not have this problem as shown in Example 6.3.

VIII. CONCLUSIONS AND FUTURE WORK

In this paper we addressed the problem of decentralized diagnosis for labeled PNs. We assume that the system is monitored by ν local sites who know the structure of the net and the initial marking, but observe its evolution with ν different masks. Each site performs diagnosis locally with a method that we previously introduced in the centralized case. We present a protocol that defines the communication rules between the coordinator and the local sites and specifies how

the diagnosis is performed by the coordinator. We proved that the proposed protocol does not produce false alarms. Moreover, we show that the absence of failure ambiguous strings is only a sufficient condition for decentralized diagnosability in the case of the considered protocol. Finally, we compare such a protocol with two other protocols we presented in (14).

One of the main goals of our future research in this topic will be that of characterizing the classes of net systems that are diagnosable in a decentralized framework using the proposed protocols even in the presence of failure ambiguous strings.

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 as close as possible to those of the centralized diagnoser.

REFERENCES

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] R. Su, W. Wonham, J. Kurien, and X. Koutsoukos, "Distributed diagnosis for qualitative systems," in *in 6th International Workshop on Discrete Event Systems, Zaragoza*, 2002, pp. 169–174.

[9] 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.

[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] 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.

[13] 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.

[14] M. Cabasino, A. Giua, A. Paoli, and C. Seatzu, "Decentralized diagnosis of Petri nets," in *Proc. 2010 American Control Conference*, 2010.

[15] C. Cassandras and S. Lafortune, *Introduction to discrete event systems, Second Edition*. Springer, 2007.

[16] M. Cabasino, A. Giua, A. Paoli, and C. Seatzu, "Decentralized diagnosability analysis of discrete event systems using Petri nets," in *Proc. 49th IEEE Conf. on Decision and Control*, 2010, submitted.