Decentralized diagnosis of discrete event systems using labeled Petri nets

Maria Paola Cabasino[1], Alessandro Giua[1], Andrea Paoli[2], Carla Seatzu[1]

[1] Department of Electrical and Electronic Engineering, University of Cagliari, Italy
e-mail: {cabasino,giua,seatzu}@diee.unica.it.

[2] Department of Electronic, Computer Science and Systems, University of Bologna, Italy
e-mail: andrea.paoli@unibo.it

Abstract

In this paper we propose an approach to the diagnosis of Petri nets in a decentralized setting that combines the decentralized scheme for automata presented by Debouk *et al.* with the diagnosis approach for Petri nets based on the notion of basis markings and justifications presented by some of the authors of this paper. The decentralized architecture that we use is composed by a set of sites communicating their diagnosis information with a coordinator that is responsible for detecting the occurrence of failures in the system. In particular, we define three 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. Finally, we prove that, as in the case of automata, diagnosability is strictly related to the presence of failure ambiguous strings.

Published as:
M.P. Cabasino, A. Giua, A. Paoli, C. Seatzu, "Decentralized diagnosis of discrete event systems using labeled Petri nets," *IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems*, Vol. PP, No. 99, pp. 1-9, Mar 2013. The original publication is available at www.ieeexplore.ieee.org.

1 Introduction

In this paper we propose an approach to the diagnosis of Petri nets (PNs) in a decentralized setting that combines the decentralized scheme for automata by Debouk *et al.* in [12] with the diagnosis approach for PNs by Cabasino *et al.* in [7, 8]. A detailed comparison between the approach presented in this paper and the approach by Debouk *et al.* is reported in the next section.

Exploiting the classical decentralized diagnosis architecture, 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 its own mask, i.e., the set of observable transitions may be different for each site. Diagnosis is locally performed using the approach founded on basis markings that we previously introduced in [7, 8]. The main feature of such an approach 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 is represented by an integer number from 0 to 3, depending on the degree of alarm. For instance, 3 is used to capture the fact that the fault has certainly occurred, whereas 0 captures the fact that the fault has not occurred. Using its own observation, each site computes the diagnosis state and, according to a given protocol, communicates it, eventually with some other information, to the coordinator who calculates global diagnosis states. In particular, three different protocols are defined that differ for the amount of information exchanged between the coordinator and the local sites, and the local sites and the coordinator. In all cases an important property is proved, namely that the coordinator never produces false alarms.

Finally, we introduce the definition of failure ambiguous strings and show that the absence of such a kind of sequences is a sufficient condition for the diagnosability of a given net system in a decentralized setting, regardless of the considered protocol. We also show that, for one of these protocols, the absence of failure ambiguous strings is also a necessary condition for the diagnosability in a decentralized setting.

The paper is organized as follows: in Section 2 a literature review is presented. In Section 3 some preliminary notions on labeled PNs are recalled, while in Section 4 the problem of decentralized diagnosis is introduced and discussed for PNs. Section 5 summarizes definitions and results on centralized diagnosis for PNs. Sections 6 and 7 contain the main results on decentralized diagnosis of PNs and on decentralized diagnosability analysis. Finally Section 8 contains conclusive remarks.

2 Literature review

In this paper we deal with the problem of decentralized fault diagnosis of discrete event systems using PNs. Solving a problem of diagnosis in the discrete event systems framework means that we associate with 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 in the centralized setting

[1, 7, 9, 11, 13, 14, 16, 17, 19, 20, 22].

Due to the intrinsically distributed nature of real systems, several distributed diagnosis techniques, that take advantage of the natural decomposition of modular systems, have been studied both in the automata [3, 10, 12, 23, 21, 24] and in the PNs setting [2, 15, 18].

In particular, focusing on PNs, in [2] Benveniste *et al.* solve a problem of alarm supervision in telecommunication networks using an unfolding approach and restricting their attention to safe PNs. In [15] Genc and Lafortune address the problem of detecting and isolating faults or other significant events in the behavior of a modular dynamic system that is modeled as a set of interacting PN modules. Faults are modeled by unobservable events and the common places capture coupling of various system components. The objective is to diagnose the occurrence of fault events based on the sequence of observed events and on the structure of the respective PN modules and their coupling by common places.

In [18] Jiroveanu and Boel propose 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 time 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 a model and from which it observes events by reconciling observations with model based predictions. The proposed algorithms use a limited information exchange between agents and can quickly ascertain whether and where a fault occurred and whether or not some components of the local processes have operated correctly. The algorithms they derive allow each local agent to generate a preliminary diagnosis prior to any communication and they show that after the communications among agents the diagnosis results are the same as in the centralized case.

Both the problem formulation and the objectives considered in [2] are significantly different from those in this paper. More strict analogies exist between our approach and the approaches of [15] and [18]. However, also in this case there exists a main difference that can be summarized as follows. In these works authors assume the PN is 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 of bordered places [15] or guard transitions [18]. 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. Thus since the problem statement is different it is not appropriate to talk about advantages or disadvantages of [18] and [15] with respect to (wrt) this paper. On the other hand, a comparison can be done with the work of Debouk *et al.* in [12] with which this work is strongly connected.

Debouk *et al.* in [12] have presented a general approach for decentralized diagnosis modeled as automata. They define three protocols – that we call D3, D2, D1 – each one characterized by a different amount of information exchanged between coordinator and local sites (the info is minimal for protocol D3 and maximal for protocol D1). Inspired by their work we consider a similar setting using PNs. Protocol D3 is very similar to our Protocol 1 because both require that a local site communicates to the coordinator only when a fault is detected. On the contrary, Protocols 2 and 3 are different from Protocols D2 and D1 because the information that coordinator and sites exchange is based on the structure of the PN. However, while in Protocol D1 each site communicates to the coordinator the corresponding state of its *extended diagnoser* and its *unobservable reach* (then an exhaustive enumeration of the possible states) for each observed event, in our Protocol 3 sites communicate to the coordinator only in some cases (when diagnosis states 2 and 3 are reached) and the information exchanged is a set of vectors (set of j-vectors). The contribution of this work is the application of a centralized diagnosis algorithm for PNs, that we propose in [7], to a decentralized setting. This requires to define the protocols on the basis of basis markings and justifications, that are the key notions of our approach. The advantages of our approach wrt the one of Debouk *et al.* is that we do not require the enumeration of the state space and we can deal with systems having an infinite state space. The disadvantage is that our approach is based on some assumptions that limit the field of applicability.

3 Background on labeled Petri nets

A *Petri 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 $M(p)$. 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 $\sigma = t_1 \ldots t_k$ 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 finite sequences that are enabled at the initial marking M_0 is denoted $L(N, M_0)$, i.e., $L(N, M_0) = \{\sigma \in T^* \mid M_0[\sigma\rangle\}$. 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 $L(N, M_0)$. Given a sequence $\sigma \in T^\star$, we call $\pi : T^\star \to \mathbb{N}^n$ the function that associates with σ a vector $y \in \mathbb{N}^n$, named *firing vector*, such that $y(t) = k$ if transition t is contained k times in σ. A firing vector y is said *minimal* if there does not exist another firing vector y' such that $\pi(y') \lneq \pi(y)$, i.e., such that each entry of y is less than or equal to the corresponding entry of y' and there exists at least one entry of y that is strictly less than the corresponding entry of y'.

A marking M is said to be *reachable* in $\langle N, M_0 \rangle$ iff there exists a 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 $R(N, M_0)$. 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 cycles is called *acyclic*. For such nets if the vector $y \in \mathbb{N}^n$ satisfies the inequality $M_0 + C \cdot y \geq \mathbf{0}$, there exists a sequence σ firable from M_0 and such that the firing vector associated with σ is equal to y. This implies that for acyclic nets $R(N, M_0) = PR(N, M_0)$.

A net system $\langle N, M_0 \rangle$ is said to be bounded if there exists a positive constant k such that for all $M \in R(N, M_0)$, $M(p) \leq k$. If such is not the case, namely if the number of tokens in one or more places can grow indefinitely, then the PN system is *unbounded*.

A *labeling function* $\mathcal{L} : T \to L \cup \{\varepsilon\}$ assigns to each transition a symbol from a given alphabet L or the empty word ε. We define $\mathcal{L}^{-1}(w) = \{\sigma \in L(N, M_0) : \mathcal{L}(\sigma) = w\}$ the *inverse operator* of $\mathcal{L}$. The set of transitions sharing the same label e is denoted T_e. Transitions whose label is ε are called *silent* or *unobservable* 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 C_u (resp. C_o) the restriction of the incidence matrix to T_u (resp. T_o). We define the *projection over* T_x, for $x \in \{u, o\}$, $P_x : T^* \to T_x^*$ as: (i) $P_x(\varepsilon) = \varepsilon$; (ii) for all $\sigma \in T^*$ and $t \in T$, $P_x(\sigma t) = P_x(\sigma)t$ if $t \in T_x$, and $P_x(\sigma t) = P_x(\sigma)$ otherwise. Given a language $K \subseteq T^*$, we denote K/σ the post-language of K after σ, i.e., $K/\sigma = \{\sigma' \in T^* \mid \sigma\sigma' \in K\}$.

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 $N' \prec_{T'} N$.

4 Problem statement

We model anomalous or faulty behavior using the set of unobservable transitions $T_f \subseteq T_u$. The set T_f includes all fault transitions and is further partitioned into r different sets T_f^i, where $i \in \mathcal{F} = \{1, \ldots, r\}$, that model different fault classes. As in most of the literature on this topic, we assume that the fault model is known, namely we know the net structure both of the fault-free and of the faulty system. The transition set $T_{reg} = T_u \setminus T_f$ represents the set of unobservable, but regular, transitions, i.e., those transitions to which a sensor is not associated but that do not describe a fault occurrence. Let

$$\bar{\mathcal{L}} : T \to L \cup \{\varepsilon\} \tag{1}$$

be the labeling function associated with the centralized system, namely the system that is able to observe all labels associated with transitions in T_o.

The problem of fault diagnosis can be seen as the problem of detecting the firing of any (unobservable) fault transition in T_f, on the basis of the observed behavior, i.e., the sequence of labels of observable transitions that have fired. 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 $\mathcal{K} = \{1, \ldots, \nu\}$ 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. Different sites have different observation masks. In particular, for each site $j \in \mathcal{K}$, the set of

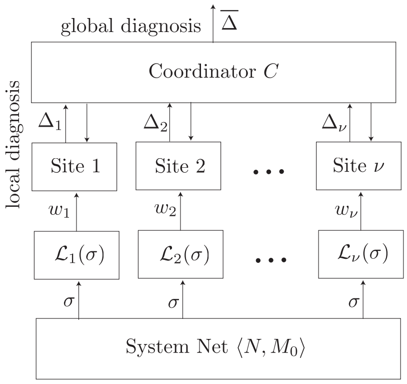

Figure 1: The decentralized diagnosis architecture.

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., $\bigcup_{j\in\mathcal{K}} T_{o,j} = T_o$. The set of locally unobservable transitions is defined as

$$T_{u,j} = T \setminus T_{o,j}. \tag{2}$$

For all $j \in \mathcal{K}$, $L_j \subseteq L$ denotes the alphabet of the j-th site, i.e., the set of labels observable by the j-th site, and

$$\mathcal{L}_j(t) = \begin{cases} \bar{\mathcal{L}}(t) & \text{if } t \in T_{o,j} \\ \varepsilon & \text{otherwise} \end{cases} \tag{3}$$

is the labeling function associated with the j-th site. Finally, $w_j = \mathcal{L}_j(\sigma)$ denotes the word of events in L_j associated with the sequence σ by the j-th site.

As shown in Fig. 1, on the basis of its own observation $w_j = \mathcal{L}_j(\sigma)$ ($j \in \mathcal{K}$) each site performs a local diagnosis. In particular, for each fault class $i \in \mathcal{F}$ it computes a different diagnosis state $\Delta_{j,i}$ (see the following Definition 3) and depending on this, it exchanges information with a *coordinator* C according to a given protocol. 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 and investigate the issue of *diagnosability* under the following assumptions.

- (**A1**) The $T_{u,j}$-induced subnet $N_{u,j}$ is acyclic for any $j \in \mathcal{K}$.
- (**A2**) The coordinator C knows which transitions can be observed by each site, i.e., it knows the sets $T_{o,j}$ for any $j \in \mathcal{K}$.
- (**A3**) 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.
- (**A4**) For each label e there exists at least one site that can observe all transitions whose label is e.

- (**A5**) Let w be a sequence of observable events generated by the PN, where such events are centrally observable. Every site must have received the projection of w (on its local alphabet) before any polling is performed by the coordinator.

Assumption A1, that is analogous to the classical hypothesis in the theory of automata where no cycle of unobservable events can appear, allows us to: (a) study the reachability of the unobservable subnet with the state equation; (b) give an easy algorithm for the computation of the firing vectors relative to justifications (see [7] for more details). Assumption A2 defines which information the coordinator has and it is necessary for the polling strategy of Protocols 2 and 3. Assumption A3 assures that the messages sent among the coordinator and the sites are not lost and are orderly received. Assumptions A4 and A5 are necessary for Protocols 2 and 3: assumption A4 guarantees the existence of a site that knows the exact number of times a given observable event has occurred; assumption A5 guarantees that the information sent and requested by the coordinator and by the local sites are relative to the same word w.

We define $\Psi(T') = \{\sigma t' \in L(N, M_0) : t' \in T'\}$, i.e., the set of all firing sequences in $L(N, M_0)$ that end with a transition $t' \in T'$. We consider the following definition of diagnosability of PNs inspired by the definition of diagnosability for (regular) languages introduced in [22].

Definition 1 *A labeled PN system $\langle N, M_0\rangle$ having no deadlock after the occurrence of any transition $t_f \in T_f^i$, for $i \in \{1, \dots, r\}$, is* diagnosable wrt the fault class T_f^i *if*

$$\begin{array}{c}\forall \sigma' \in \Psi(T_f^i), \quad \exists K \in \mathbb{N}, \quad \forall \sigma'' \in L(N, M_0)/\sigma', \\ |\sigma''| \geq K \;\Rightarrow\; \forall \sigma \in \mathcal{L}^{-1}(\mathcal{L}(\sigma'\sigma'')),\; \exists t_f \in T_f^i \;:\; t_f \in \sigma\end{array} \tag{4}$$

A labeled PN system $\langle N, M_0\rangle$ is said to be diagnosable *if it is diagnosable wrt all fault classes.* ■

In words, given a firing sequence σ' that ends in a fault transition, let σ'' be any sufficiently long continuation of it, i.e., $|\sigma''| \geq K$, where K depends on σ'. A labeled PN system $\langle N, M_0\rangle$ having no deadlock after the occurrence of any transition $t_f \in T_f^i$, for $i \in \{1, \dots, r\}$, is *diagnosable wrt the fault class T_f^i* if any firing sequence σ belonging to the language and having the same observable projection of $\sigma'\sigma''$ contains a fault transition in T_f^i. This implies that along any continuation σ'' of σ' the occurrence of a fault transition in T_f^i can be detected in a finite number of transitions firings (at most K).

5 Basic definitions and results on centralized diagnosis

In this section we briefly recall the diagnosis procedure we defined in [7, 8] in the centralized setting, 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 with transitions in T_o. In particular, we first provide some preliminary definitions.

- Let $w = \mathcal{L}(\sigma)$ be the word of events associated with 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.

- Given a word $w \in L^*$, let $\sigma_o \in T_o^*$ be a sequence of observable transitions such that $\mathcal{L}(\sigma_o) = w$. A *basis marking* M_b is a marking reached from M_0 with the firing of σ_o and of all unobservable transitions whose firing is *strictly* necessary to enable w. Such a sequence σ_u of unobservable transitions interleaved with σ_o whose firing enables σ_o and whose firing vector is *minimal* is called *justification*. Since in general σ_o is not unique and more than one σ_u may be associated with each σ_o, then the set of justifications of w is not a singleton.

- We denote
$$\begin{aligned}
&\hat{\mathcal{J}}(w) = \\
&\quad \{\ (\sigma_o, \sigma_u),\ \sigma_o \in T_o^*,\ \mathcal{L}(\sigma_o) = w,\ \sigma_u \in T_u^*\ \mid \\
&\quad [\exists \sigma \in \mathcal{S}(w)\ :\ \sigma_o = P_o(\sigma),\ \sigma_u = P_u(\sigma)] \land \\
&\quad [\nexists \sigma' \in \mathcal{S}(w) :\ \sigma_o = P_o(\sigma'), \\
&\quad\ \ \sigma'_u = P_u(\sigma')\ \land\ \pi(\sigma'_u) \lneq \pi(\sigma_u)]\}
\end{aligned}$$
the set of pairs (sequence $\sigma_o \in T_o^*$ with $\mathcal{L}(\sigma_o) = w$ - corresponding *justification* of w). Let
$$\begin{aligned}
&\hat{Y}_{\min}(M_0, w) = \\
&\quad \{(\sigma_o, y), \sigma_o \in T_o^*, \mathcal{L}(\sigma_o) = w, y \in \mathbb{N}^n \mid \\
&\qquad\qquad \exists (\sigma_o, \sigma_u) \in \hat{\mathcal{J}}(w) : \pi(\sigma_u) = y\}
\end{aligned}$$
be the set of pairs (sequence $\sigma_o \in T_o^*$ with $\mathcal{L}(\sigma_o) = w$, corresponding *j-vector*). In simple words, $\hat{\mathcal{J}}(w)$ is the set of pairs whose first element is the sequence $\sigma_o \in T_o^*$ labeled w and whose second element is the firing vector of the corresponding sequence of unobservable transitions interleaved with σ_o whose firing enables σ_o and whose firing vector is minimal. The firing vectors of these sequences are called *j-vectors*. Finally, let us denote Y the set of j-vectors for the observed word w.

Example 2 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. The set of sequences consistent with the actual observation is $\mathcal{S}(w) = \{\varepsilon_4 t_1 t_2, \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3, \varepsilon_4 t_1 t_2 \varepsilon_4, \varepsilon_4 t_1 t_2 \varepsilon_5, \varepsilon_4 t_1 t_2 \varepsilon_5 \varepsilon_6, \varepsilon_4 t_1 t_2 \varepsilon_5 \varepsilon_6 \varepsilon_7, \varepsilon_4 t_1 t_2 \varepsilon_5 \varepsilon_6 \varepsilon_7 \varepsilon_8, \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3 \varepsilon_4, \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8$ $\varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3 \varepsilon_5 \varepsilon_6, \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3 \varepsilon_5 \varepsilon_6 \varepsilon_7, \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 t_3 \varepsilon_5 \varepsilon_6 \varepsilon_7 \varepsilon_8\}$. The set of pairs (sequence $\sigma_o \in T_o^*$ with $\mathcal{L}(\sigma_o) = w$ - corresponding justification of w) is $\hat{\mathcal{J}}(w) = \{(t_1 t_2, \sigma_1), (t_1 t_3, \sigma_2)\} = \{(t_1 t_2, \varepsilon_4), (t_1 t_3, \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8)\}$ Note that, $\sigma_3 = \varepsilon_4 \varepsilon_4$, $\sigma_4 = \varepsilon_4 \varepsilon_5$, $\sigma_5 = \varepsilon_4 \varepsilon_5 \varepsilon_6$, $\sigma_6 = \varepsilon_4 \varepsilon_5 \varepsilon_6 \varepsilon_7$, $\sigma_7 = \varepsilon_4 \varepsilon_5 \varepsilon_6 \varepsilon_7 \varepsilon_8$, $\sigma_8 = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8 \varepsilon_4$, $\sigma_9 = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8 \varepsilon_5$, $\sigma_{10} = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8 \varepsilon_5 \varepsilon_6$, $\sigma_{11} = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8 \varepsilon_5 \varepsilon_6 \varepsilon_7$ and $\sigma_{12} = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8 \varepsilon_5 \varepsilon_6 \varepsilon_7 \varepsilon_8$ are not

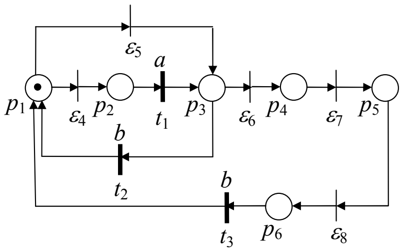

Figure 2: The PN system considered in Examples 2 and 4.

justifications since their firing vector is not minimal. As an example,

$$\pi(\sigma_1) = \begin{array}{c} \left[\begin{array}{cccccccc} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array}\right]^T \\ \begin{array}{cccccccc} t_1 & t_2 & t_3 & \varepsilon_4 & \varepsilon_5 & \varepsilon_6 & \varepsilon_7 & \varepsilon_8 \end{array} \end{array} \lneq$$

$$\pi(\sigma_3) = \begin{array}{c} \left[\begin{array}{cccccccc} 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 \end{array}\right]^T \\ \begin{array}{cccccccc} t_1 & t_2 & t_3 & \varepsilon_4 & \varepsilon_5 & \varepsilon_6 & \varepsilon_7 & \varepsilon_8 \end{array} \end{array}.$$

The set of pairs j-vectors (sequence $\sigma_o \in T_o^*$ with $\mathcal{L}(\sigma_o) = w$, corresponding j-vector) is $\hat{Y}_{min}(M_0, w) = \{(t_1 t_2, [1\ 0\ 0\ 0\ 0]^T), (t_1 t_3, [1\ 0\ 1\ 1\ 1]^T)\}$ and they all lead to the same basis marking $M_0 = [2\ 0\ 0\ 0\ 0\ 0]^T$. ■

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

Definition 3 *A* diagnoser *is a function $\Delta : L^* \times \{T_f^1, T_f^2, \ldots, T_f^r\} \rightarrow \{0, 1, 2, 3\}$ that associates with each observation w and each fault class T_f^i, $i = 1, \ldots, r$, 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 $t_f \not\in \sigma$.*
  In such a case the i-th fault cannot have occurred, because none of the sequences consistent with the observation contains fault transitions in T_f^i.
- $\Delta(w, T_f^i) = 1$ *if:*
  1. *there exist $\sigma \in \mathcal{S}(w)$ and $t_f \in T_f^i$ such that $t_f \in \sigma$ but*
  2. *for all $(\sigma_o, \sigma_u) \in \hat{\mathcal{J}}(w)$ and for all $t_f \in T_f^i$ it holds that $t_f \not\in \sigma_u$.*

  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*
  1. *there exists $t_f \in T_f^i$ such that $t_f \in \sigma_u$;*
  2. *for all $t_f \in T_f^i$, $t_f \not\in \sigma_u'$.*

In such a case a fault transition in the i-th class is contained in at least 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. ■

Note that we associate a diagnosis state equal to 1 when the fault may have occurred but it is not contained in any justification of the considered word, while we associate a diagnosis state equal to 2 when the fault is contained in at least one (but not all) justification of the considered word. A systematic procedure has been given in [7, 8] to compute the above diagnosis states that is not recalled here for the sake of brevity.

Example 4 Let us consider again the PN in Fig. 2, where $T_f = \{\varepsilon_5, \varepsilon_7\}$. Let $w = a$. In such a case it is $\Delta(w, T_f) = 1$. In fact, $\hat{\mathcal{J}}(a) = \{(t_1, \varepsilon_4)\}$ but there exists $\sigma = \varepsilon_4 t_1 \varepsilon_6 \varepsilon_7 \varepsilon_8 \in \mathcal{S}(a)$ containing the fault ε_7. Finally, let $w = ab$. In such a case it is $\Delta(w, T_f) = 2$. In fact, as shown in Example 2, the justifications of ab are: $\sigma_1 = \varepsilon_4$, that does not contain fault transitions and $\sigma_2 = \varepsilon_4 \varepsilon_6 \varepsilon_7 \varepsilon_8$ that contains $\varepsilon_7 \in T_f$. ■

6 Decentralized diagnosis

In this section we introduce three different protocols to solve the decentralized diagnosis problem introduced in Section 4 [4, 5]. In the following we denote 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.

6.1 Diagnosis under Protocol 1

Protocol 1 is based on the following very simple rules illustrated in Algorithm 5.

Algorithm 5 [Algorithm for Protocol 1]

```
1. Each site $j \in \mathcal{K}$:
  1.a. sets $w_j = \varepsilon$;
  1.b. computes its diagnosis state $\Delta(w_j, T_f^i)$ for all $i \in \mathcal{F}$.
2. The diagnosis state of the coordinator $\bar{\Delta}_i$ relative to
   each $T_f^i$, for all $i \in \mathcal{F}$, is initially undefined.
3. Wait until a new transition $t \in T_o$ fires.
4. Each site $j \in \mathcal{K}$:
  4.a. sets $w'_j = w_j$ and $w_j = w'_j \mathcal{L}_j(t)$;
  4.b. computes its diagnosis state $\Delta_{j,i} = \Delta(w_j, T_f^i)$
       for all $i \in \mathcal{F}$.
```

4.c. If $\Delta_{j,i} = 3$ and $\Delta_{j,i} > \Delta(w'_j, T_f^i)$ for some $i \in \mathcal{F}$,
then transmits to the coordinator its diagnosis state.

5. If the coordinator receives a diagnosis state $\Delta_{j,i} = 3$
from any site $j \in \mathcal{K}$, it sets $\bar{\Delta}_i = 3$ (fault).

6. Go to step 3.

■

A decentralized diagnoser using Protocol 1 satisfies the following important property.

Proposition 6 *Under assumptions A1 and A3 the coordinator based on Protocol 1 never produces false alarms, namely if $\bar{\Delta}_i = 3$, then $\Delta_i^* = 3$ as well.*

Proof By assumption A3, if the coordinator diagnosis state is $\bar{\Delta}_i = 3$, it means that there exists at least one site $j \in \mathcal{K}$ such that $\Delta_{j,i} = 3$. Now, by eq. (2) 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. Note that assumption A1 is necessary for the computation of the justification (see [7]). □

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 7 Let us consider the PN 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 transitions are $T_{o,1} = \{t_1, t_4, t_5\}$ and $T_{o,2} = \{t_2, t_3, t_5\}$, respectively. Thus, the 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 a unique 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 as well.

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

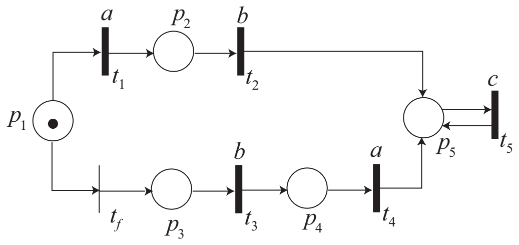

Figure 3: The PN system considered in Example 7.

Let us now discuss diagnosability. From Proposition 6 the following result obviously holds.

Corollary 8 *If a system is diagnosable in the decentralized setting (regardless of the used protocol), then it is also diagnosable in the centralized setting.* ■

Clearly, the reverse of the implication does not hold. However, in the case of diagnosis performed using Protocol 1 the following result can be proved. We address to Theorem 11 in Section 6.4 of [12] for the proof of this result.

Proposition 9 *The system is diagnosable wrt the decentralized approach based on Protocol 1 iff for every fault class $i \in \mathcal{F}$ there exists at least one site $j \in \mathcal{K}$ such that the system is diagnosable by the j-th site wrt that fault class.*

6.2 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 local sites to recompute their diagnosis states. This may lead to an improvement of the quality of the diagnosis achieved by the coordinator.

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

- Let $\mathcal{K}_e = \{k \in \mathcal{K} \mid \forall t \in T : \mathcal{L}(t) = e \Rightarrow t \in T_{k,o}\}$ be the set of sites (by assumption A4 this set is never empty) that are capable of observing all transitions labeled e.
- Given a site j and a set of j-vectors Y_j,

$$\begin{aligned}\mathcal{I}(j, Y_j) = \ & \{e \in L \ \mid \exists\, y \in Y_j \ \wedge \exists\, t \in T \setminus T_{o,j} \ : \\ & \qquad\qquad\qquad y(t) > 0 \wedge \mathcal{L}(t) = e\}\end{aligned}$$

  is the set of labels relative to transitions that appear in at least a j-vector of the j-th module.
- Let $|w_k|_e$ be the number of occurrences of label e in the observation w_k.

- Given an observation w_k from site k, a label e, and a j-vector y,

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

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

Based on the above definitions, we can summarize the main steps of the decentralized procedure based on Protocol 2 with the following algorithm. The idea beyond the algorithm is that some justifications of a site transmitted to the coordinator can be confuted with the knowledge of the information of other sites. In particular, let consider the refinement of Y_j. If Y_j contains a j-vector that assumes a certain number of occurrences of e, but this number is not consistent with the observation of a site that is capable of observing e, then such a justification is certainly unfeasible. Therefore, if $\beta_k(w_k, e, y) < 0$ for a certain label e 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 the j-vector y assumes a number of occurrences of e that is greater than the real number, that is exactly known by the k-th site. On the contrary, if $\beta_k(w_k, e, y) \geq 0$ it means that the j-vector y is compatible with the observation of the k-th site. In particular, if $\beta_k(w_k, e, y) = 0$ it means that the justification contains the same number of occurrences of label e as those observed by site k. If $\beta_k(w_k, e, y) > 0$ it means that the justification relative to y does not contain all the occurrences of e; thus the rest of transitions labeled e, up to the value $|w_k|_e$, have fired after the justification and the observation w_j. Finally, in the formulation of the algorithm we assume that a new transition cannot fire until the procedure of communication and polling among the coordinator and the sites is ended.

Algorithm 10 [Algorithm for Protocol 2]

```
1. Each site $j \in \mathcal{K}$:
   1.a. sets $w_j = \varepsilon$;
   1.b. computes its diagnosis state $\Delta(w_j, T_f^i)$ for all $i \in \mathcal{F}$.
2. The diagnosis state $\bar{\Delta}_i$ of the coordinator relative to each $T_f^i$, for all $i \in \mathcal{F}$, is initially undefined.
3. Wait until a new transition $t \in T_o$ fires.
4. Each site $j \in \mathcal{K}$:
   4.a. sets $w'_j = w_j$ and $w_j = w'_j \mathcal{L}_j(t)$;
   4.b. computes its diagnosis state $\Delta_{j,i} = \Delta(w_j, T_f^i)$ for all $i \in \mathcal{F}$.
   4.c. If $\Delta_{j,i} = 3$ and $\Delta_{j,i} > \Delta(w'_j, T_f^i)$ for some $i \in \mathcal{F}$, then transmits to the coordinator its diagnosis state and its set of j-vectors $Y_j$.
5. Let $\mathcal{K}' \subseteq \mathcal{K}$ be the set of all sites that have transmitted their diagnosis states to the coordinator in step 4.c.
   For all $i \in \mathcal{F}$ the coordinator sets $\bar{\Delta}_i = 3$ if at step 4.c it has received a diagnosis state $\Delta_{j,i} = 3$
```

```
  from some $j \in \mathcal{K}'$.
6. Let $\mathcal{W}$ be a row vector having as many entries as the number of labels in $\bar{\mathcal{L}}$ and let initially set $Nan$ each entry in $\mathcal{W}$.
7. For each site $j \in \mathcal{K}'$:
  7.a. the coordinator computes $\mathcal{I}(j, Y_j)$;
  7.b. for each label $e \in \mathcal{I}(j, Y_j)$
        7.b.i. If the entry of $\mathcal{W}$ corresponding to $e$ is equal to $Nan$, the coordinator polls one site $k \in \mathcal{K}_e$ to know the value of $|w_k|_e$ and stores this number in the corresponding entry of $\mathcal{W}$;
        7.b.ii. If $\beta_k(w_k, e, y) < 0$ 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.
        7.b.iii. 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.b.iv. 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 them to the coordinator that sets the corresponding $\bar{\Delta}_i$ equal to 3.
8. Go to step 3.
```

■

Note that the vector $\mathcal{W}$ ensures that for any label $e \in \mathcal{I}(j, Y_j)$ no more than one polling is done for a given sequence of transitions firing.

The refinement process on which Protocol 2 is based has in general positive effects on diagnosis as shown by the following example.

Example 11 Let us consider the PN 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 $T_{o,1} = \{t_3, t_4, t_9\}$ and $T_{o,2} = \{t_1, t_2, t_5, t_6\}$, respectively. Assume that $L_1 = \{a, c\}$ and $L_2 = \{b\}$, thus $\mathcal{K}_a = \{1\}$, $\mathcal{K}_b = \{2\}$ and $\mathcal{K}_c = \{1\}$.

If no transition fires we have $\sigma = \varepsilon$, thus $w = \varepsilon$. For the first site $\Delta_{1,1} = 1$, relative to the first fault class, $\Delta_{1,2} = 1$, relative to the second fault class, because at initial marking the sequence $t_{f,1}^I t_1 t_{f,2}$ of unobservable transitions may have fired. On the other hand, $\Delta_{2,1} = 1$ and $\Delta_{2,2} = 0$, because no sequence of unobservable transitions enabled at the initial marking contains the fault transition $t_{f,2}$. Thus no site communicates to the coordinator.

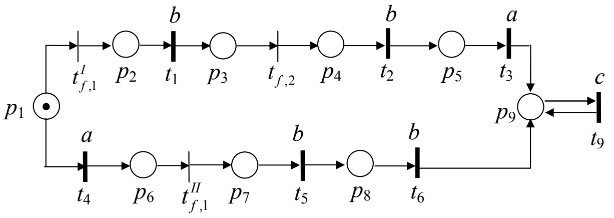

Figure 4: The PN system considered in Example 11.

Now let $\sigma = t^I_{f,1}t_1$, thus $w = b$. The first site has $w_1 = \varepsilon$, while the second site has $w_2 = b$. Then, as at the previous step $\Delta_{1,1} = \Delta_{1,2} = 1$, while for the second site $\Delta_{2,1} = 3$ and $\Delta_{2,2} = 1$. Site 2 transmits $\Delta_{2,1} = 3$ to the coordinator together with its set of j-vectors: $Y_2 = \{y_2', y_2''\}$, where y_2' is the firing vector relative to $\sigma'_{u,2} = t^I_{f,1}$, while y_2'' is the firing vector relative to $\sigma''_{u,2} = t_4 t^{II}_{f,1}$.

Since $\mathcal{I}(2, Y_2) = \{a\}$ and $\mathcal{K}_a = \{1\}$, the coordinator polls site 1 to know the number of symbols a it has observed. Since $|w_1|_a = 0$, then $\beta_1(w_1, a, y_2'') = 0 - 1 < 0$. It means that j-vector y_2'' can be confuted and removed from Y_2. The refined set of j-vectors is $Y_2' = \{y_2'\}$ and it is communicated to site 2. There is no change in the diagnosis state, however this refinement will allow site 2 to detect the fault of the second fault class at the next observation.

Finally let $\sigma = t^I_{f,1}t_1 t_{f,2} t_2$, thus $w = bb$. It is $w_1 = \varepsilon$ and $w_2 = bb$. Then, again $\Delta_{1,1} = \Delta_{1,2} = 1$, while for the second site $\Delta_{2,1} = 3$ and $\Delta_{2,2} = 3$. Site 2 transmits $\Delta_{2,1} = \Delta_{2,2} = 3$ to the coordinator together with its set of j-vectors: $Y_2 = \{y_2'''\}$, where y_2''' is the firing vector relative to $\sigma'''_{u,2} = t^I_{f,1}t_{f,2}$. Since now $\mathcal{I}(2, Y_2) = \emptyset$ the coordinator does not start the polling procedure.

Note that, the firing of one transition in T^1_f is detected using both 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 thanks to the refining procedure of the set of j-vectors through the polling of the coordinator.

■

Remark 12 Let us now discuss the effects of delays in Protocol 2.

Since events occur in an asynchronous way, i.e., we are not assuming that there is a global clock, it can obviously happen that the value of $|w_k|_e$, i.e., the number of occurrences of label e in the observation w_k, which the coordinator requests from the polled sites, is affected by some delay. As a result of this the coordinator may receive a value $|w_k|'_e > |w_k|_e$ because during the delay between the start of the polling and the arrival of the request to the kth polled site other transitions labeled e may have fired. This implies that the the difference between the number of times the site k has observed e and the number of times a transition labeled e appears in y, namely $\beta_k(w_k, e, y)$, may be greater than the correct one. In particular, it may occur that a negative value of $\beta_k(w_k, e, y)$ 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. As an example let us consider Example 11 when $w = b$ is observed. It could happen that while site 2 communicates its diagnosis state to the coordinator, transitions $t_{f,2}t_2t_3$ fire. In such a case when the coordinator polls site 1 to know how many a's it has observed, namely to know $|w_1|_a$, site 1 answers $|w_1|_a = 1$, because its new observation is now $w_2 = a$. Thus if no delay occurs (as in the case considered in Example 11) $\beta_1(w_1, a, y_2'') = 0 - 1 < 0$ then we can reject j-vector y_2'' and detect the occurrence of $t_{f,2}$ at the next observation. If a delay occurs, it may happen that the advantages of Protocol 2 are lost, but in any case no false alarm can occur. In the considered example if the described delay is considered $\beta_1(w_1, a, y_2'') = 1 - 1 = 0$ (because w_1 has changed).

Due to the absence of a global clock it may also happen that after the polling, the coordinator transmits the refined set of j-vectors Y_j' to site j, but in the meanwhile site j has observed another event and has computed the diagnosis state on the basis of the old and not refined set Y_j. Also in this case such a delay may never cause false alarms, but only avoid the occurrence of the refinement that leads to a better estimation. To better understand, let us consider again Example 11 when $w = b$ is observed. It could happen that in the meanwhile that the coordinator polls site 1, transitions $t_{f,2}t_2$ fire. In such a case site 2 computes its new set of j-vectors and its new diagnosis state on the basis of $Y_2 = \{y_2', y_2''\}$. Thus when the coordinator will communicate to site 2 the refined set of j-vectors $Y_2' = \{y_2'\}$ it cannot use this information anymore. Even in this case if a delay occurs, it may happen that the advantages of Protocol 2 are lost. ■

The following propositions can be stated.

Proposition 13 *Under assumptions A1 to A5 the coordinator based on Protocol 2 never produces false alarms, namely if $\bar{\Delta}_i = 3$, then $\Delta_i^* = 3$ as well.*

Proof By Proposition 6 (where assumptions A1 and A3 must hold) we know that no false alarm may occur when using Protocol 1. Now, by assumptions A2, A4 and A5 the effect of Protocol 2 is that of eventually reducing the cardinality of the sets of j-vectors relative to certain sites, wrt those computed using Protocol 1. In fact, the coordinator knows which sites should be polled (assumptions A2 and A3) to know the exact number of times a given observable event e has occurred. This number, in turn, is an upper bound on the number of times that event e has occurred in a feasible justification (assumption A4). By definition such a reduction consists in only removing those j-vectors that are certainly not feasible, because they are not consistent with the observations of other sites. Finally, assumption A5 guarantees that all sites and the coordinator are referring to the same word w. Thus Protocol 2 guarantees that no false alarm may occur as well. □

Proposition 14 *The 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 $j \in \mathcal{K}$ that perform a refinement of $Y_{min}(M_0, w_j)$.*

Proof The result follows from the fact that the set $Y_{min}(M_0, w_j)$ certainly contains 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 may be confuted, but certainly vector $\bar{y}$ will not, therefore $\bar{y} \in Y'_{min}(M_0, w_j)$, thus proving the statement. □

Proposition 15 *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 $j \in \mathcal{K}$ such that the system is diagnosable by the j-th site wrt that fault class.*

Proof For simplicity, with no loss of generality, we assume that there is only one fault class. If there exists one site $j \in \mathcal{K}$ such that the system is diagnosable by the j-th site, due to Assumption **A1**, this means that the j-th site certainly reconstructs 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. □

The above proposition only provides a sufficient condition for diagnosability. In fact, it may happen that the system is locally not diagnosable by any site, while it is diagnosable in a decentralized setting.

This is the case of the PN system in Example 11. In fact, both diagnosers of the systems observing $T_{o,1} = \{t_3, t_4, t_9\}$ and $T_{o,2} = \{t_1, t_2, t_5, t_6\}$ are not able to detect the occurrence of $t_{f,2}$ if the sequence $\sigma = t_{f,1}^{I} t_1 t_{f,2} t_2 t_3 t_6^k$ fires, where k is an arbitrary integer number. On the contrary, as shown in Example 11, 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.

Example 16 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 sets of transitions are $T_{o,1} = \{t_1, t_2, t_3, t_6\}$ and $T_{o,2} = \{t_4, t_5, t_6\}$ and 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 $Y_{min}(M_0, w_1) = Y_1 = \{y_1', y_1''\}$ 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 $Y_{min}(M_0, w_2) = Y_2 = \{y_2', y_2''\}$ 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 setting, 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 $\sigma_u = t_f$. ■

6.3 Diagnosis under Protocol 3

Protocol 3 differs from Protocol 2 for the fact that each site communicates its diagnosis state and its set of j-vectors to the coordinator, not only when the diagnosis state is equal to 3, but also when it is equal to 2.

Thus the main steps of the decentralized procedure based on Protocol 3 are the same as those

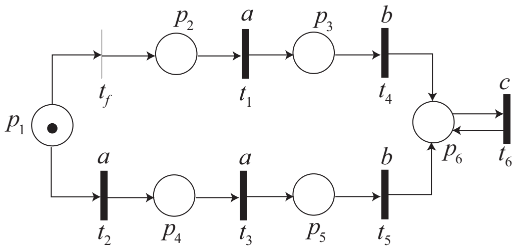

Figure 5: The PN system considered in Example 16.

relative to Protocol 2 apart from the fact that in Step 5 of Algorithm 10 the sites in $\mathcal{K}'$ are those whose $\Delta_{j,i} = \{2, 3\}$ rather than $\Delta_{j,i} = 3$.

As intuitive, a greater number of information exchanged leads to better diagnosis capability as shown by the following example.

Example 17 Let us consider again the PN in Fig. 5 where $T_u = T_f = \{t_f\}$. The net is locally diagnosed by two sites whose sets of observable transitions are $T_{o,1} = \{t_1, t_2, t_3, t_6\}$ and $T_{o,2} = \{t_4, t_5, t_6\}$, respectively. This implies that $L_1 = \{a, c\}$, $L_2 = \{b, c\}$, $\mathcal{K}_a = \{1\}$, $\mathcal{K}_b = \{2\}$ and $\mathcal{K}_c = \{1, 2\}$. Let us 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 for the first site is $Y_{min}(M_0, w_1) = Y_1 = \{y_1', y_1''\}$, where $y_1' = \vec{0}$ and $y_1'' = \pi(t_f)$, while for the second site is $Y_{min}(M_0, w_2) = Y_2 = \{y_2', y_2''\}$, where $y_2' = \pi(t_f t_1)$ 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 $\mathcal{K}_a = \{1\}$. Thus the coordinator polls site 1 to know the number of a labels 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 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 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 t_f is detected. ■

The following important property can also be demonstrated in the case of Protocol 3.

Proposition 18 *Under assumptions A1 to A5 the coordinator under Protocol 3 does not produce any false alarm, namely if $\bar{\Delta}_i = 3$, then $\Delta_i^* = 3$ as well.*

Proof It can be proved following the same arguments of Proposition 13. □

We conclude this section with a remark.

Clearly, several other protocols can be defined. The choice of the most appropriate protocol corresponds to the determination of the best trade-off between the amount of information exchanged and the diagnosis capabilities, that obviously depends on the particular application. If we want a protocol that has the same performances of the centralized diagnoser we need to synchronize at each step and to ask all sites at each step to send all the consistent states to the coordinator. Then the coordinator does an intersection of all consistent states of all sites and obtains the same information that the centralized diagnoser has. Another possibility to obtain the same performance of the centralized diagnoser is to increase the knowledge of the coordinator: as an example if the coordinator knows the structure of the net and what the different sites can observe, each site can just send its own observation and the coordinator computes with these information the set of consistent markings of each site and does the intersection; another case is when the coordinator knows the unobservable reach of each site, each site can send the set of basis markings and justifications at each step and the coordinator computes with these information the set of consistent markings of each site and does the intersection.

7 Diagnosability and failure ambiguous strings

In this section we introduce the definition of failure ambiguous strings, and show the relationships among them. We want to show that, regardless of the used protocol, when analyzing diagnosability in a decentralized setting, the first important step is that of detecting the presence of particular strings, called *failure ambiguous strings.*

Note that the notion of failure ambiguous strings has been firstly introduced in [12] in the setting of automata under the assumption of two sites. Here we extend such definition to PNs and consider the general case of an arbitrary number ν of sites.

Definition 19 *Consider a net system $\langle N, M_0 \rangle$ monitored by a set $\mathcal{K} = \{1, \ldots, \nu\}$ of sites. Let $T_{o,j} \subseteq T_o$ be the set of locally observable transitions for site $j \in \mathcal{K}$. Finally, let $T_f^i \subseteq T_f$ be the i-th fault class, with $i \in \mathcal{F}$.*

A string $\sigma \in T^*$ 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 hold:*

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

(b) $\bar{\mathcal{L}}^{-1}(\bar{\mathcal{L}}(\sigma)) \cap (T \setminus T_f^i)^* = \emptyset$,

where $\bar{\mathcal{L}}$ and $\mathcal{L}_j$ are defined as in (1) *and* (3)*, respectively.*

■

In simple words, a sequence σ is failure ambiguous wrt to a set of sites and the ith fault class if the following conditions are simultaneously verified: 1) for all sites $j \in \mathcal{K}$ the word σ is uncertain, i.e., produces an uncertain diagnosis state $\Delta_{j,i} \in \{1, 2\}$; and 2) the word σ is not uncertain for the centralized system.

Example 20 Let us consider the PN system in Fig. 6 which is locally diagnosed by two sites whose sets of transitions are $T_{o,1} = \{t_1, t_3, t_5, t_6, t_7\}$ and $T_{o,2} = \{t_2, t_3, t_4, t_5, t_7\}$, and whose alphabets are equal to $L_1 = \{a, c\}$ and $L_2 = \{b, c\}$, respectively.

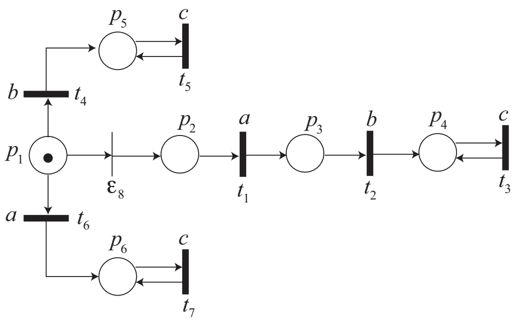

Figure 6: Petri net system for Example 20.

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 $\bar{\mathcal{L}}(\sigma) = \{abc^q\}$ and $\bar{\mathcal{L}}^{-1}(\bar{\mathcal{L}}(\sigma)) = \{\varepsilon_8 t_1 t_2 t_3^q\}$ thus $\bar{\mathcal{L}}^{-1}(\bar{\mathcal{L}}(\sigma)) \cap (T \setminus T_f^i)^* = \emptyset$. ■

In general cases, as it happens in the case of automata [12], the absence of failure ambiguous strings of arbitrary length is only a sufficient condition for the diagnosability in a decentralized setting. In fact, if protocols are defined so that local sites take advantage of the information collected by the other sites (e.g., receiving certain information by the coordinator), the resulting system may be diagnosable even in the presence of failure ambiguous strings. On the contrary, if each site computes its diagnosis states receiving no information from the other sites and from the coordinator, then the absence of failure ambiguous strings is also a necessary condition for the decentralized diagnosability.

Using Protocol 1, where a site communicates to the coordinator its diagnosis state if and only if it has detected the occurrence of a fault and no communication is allowed among sites, and from the coordinator to the local sites, it is obvious that the absence of failure ambiguous strings arbitrarily long after the fault is not only a sufficient condition for decentralized diagnosability, but it is also necessary. On the contrary, if we use the more sophisticated protocols, as Protocol 2 and 3, it may occur that a system is diagnosable in a decentralized setting even in the presence of failure ambiguous strings. This is due to the fact that the protocol is based on a confutation procedure that allows the sites to take benefit of the information sent by the other sites to the coordinator.

Example 21 Let us consider the PN system in Fig. 7 where $T_u = T_f = \{\varepsilon_{10}\}$. The net is monitored by two sites whose set of observable transitions is respectively $T_{o,1} = \{t_1, t_2, t_3, t_6, t_9\}$ and $T_{o,2} = \{t_4, t_5, t_6, t_7, t_8\}$. This implies that $L_1 = \{a, c\}$, $L_2 = \{b, c\}$, $\mathcal{K}_a = \{1\}, \mathcal{K}_b = \{2\}$ and $\mathcal{K}_c = \{1, 2\}$.

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 $\mathcal{L}_2^{-1}(\mathcal{L}_2(\sigma)) \cap (T \setminus T_f)^* =$

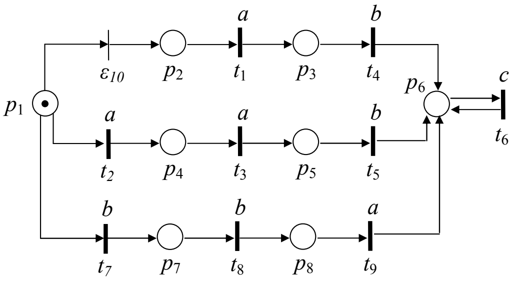

Figure 7: The Petri net system considered in Example 21.

$\{t_2t_3t_5t_6^q\}$; and $\bar{\mathcal{L}}(\sigma) = \{abc^q\}$ and $\bar{\mathcal{L}}^{-1}(\bar{\mathcal{L}}(\sigma)) = \{\varepsilon_{10}t_1t_4t_6^q\}$ thus $\bar{\mathcal{L}}^{-1}(\bar{\mathcal{L}}(\sigma)) \cap (T \setminus T_f^i)^* = \emptyset$.

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 $Y_{\min}(M_0, w_1) = \{\pi(\varepsilon_{10}t_4), \pi(t_7t_8)\}$ and $\Delta_1 = 2$. Similarly, the second site observes $w_2 = bc^q$ thus $Y_{\min}(M_0, w_2) = \{\pi(\varepsilon_{10}t_1), \pi(t_2t_3)\}$ and $\Delta_2 = 2$ as well. However, both $\pi(t_7t_8)$ and $\pi(t_2t_3)$ 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 of arbitrary length, we can conclude that the system is diagnosable using Protocol 3 even in the presence of failure ambiguous strings of arbitrary length after the fault being the centralized system diagnosable. ■

Obviously, regardless of the considered protocol, if a system is diagnosable in a centralized setting wrt a given fault class, and has no failure ambiguous string of arbitrary length wrt that class, it is also diagnosable in a decentralized setting. The following proposition formally proves this.

Proposition 22 *Consider a net system $\langle N, M_0 \rangle$ monitored by a set $\mathcal{K} = \{1, \ldots, \nu\}$ of sites. Let $T_f^i \subseteq T_f$ be the generic i-th fault class, with $i \in \mathcal{F}$. Let us suppose that the net system $\langle N, M_0 \rangle$ is diagnosable in a centralized setting wrt T_f^i.*

If there do not exist failure ambiguous strings of arbitrary length for the considered set of sites wrt to T_f^i, then the system is also diagnosable in a decentralized setting using Protocol 1, 2 or 3 to perform decentralized diagnosis.

Proof By Definition 19, if there do not exist failure ambiguous strings of arbitrary length wrt a given fault class, it means that there do not exist strings of arbitrary length that can be distinguished by the centralized diagnoser, but cannot be distinguished by all the local sites. This implies that, for each string containing a fault there exists at least one site that detects the fault. Thus if the system is diagnosable in a centralized setting, then it is also diagnosable in a decentralized setting. □

In [6] we have presented a procedure to verify the absence of such kind of strings for both bounded and unbounded PN systems.

8 Conclusions and future work

The contribution of this paper consists in the definition of three protocols for the decentralized diagnosis of discrete event systems using labeled PNs. It is proved that all such protocols prevent false alarms, while their diagnosability properties depend on the amount of information exchanged with a central coordinator.

Several lines of investigations remain to be explored, including: (i) relaxation of some assumptions that characterize our decentralized diagnosis approach; (ii) characterization of the effects that delays have on our procedure; (iii) consideration of the case where the coordinator always produces a diagnosis state, that may also be an uncertain or a non faulty state: in such a case appropriate protocols should be defined assuming information exchanges among the local sites and the coordinator also in the case of local diagnosis states equal to 0 and 1.

References

[1] F. Basile, P. Chiacchio, and G. De Tommasi. An efficient approach for online diagnosis of discrete event systems. *IEEE Trans. Automatic Control*, 54(4):748–759, 2009.

[2] A. Benveniste, E. Fabre, S. Haar, and C. Jard. Diagnosis of asynchronous discrete event systems, a net unfolding approach. *IEEE Trans. Automatic Control*, 48(5):714–727, 2003.

[3] R.K. Boel and J.H. van Schuppen. Decentralized failure diagnosis for discrete-event systems with costly communication between diagnosers. In *Proc. 6th Work. on Discrete Event Systems (Zaragoza, Spain)*, October 2002.

[4] M. Cabasino, A. Giua, A. Paoli, and C. Seatzu. A new protocol for the decentralized diagnosis of labeled Petri nets. In *Proc. 10th Work. on Discrete Event Systems (Berlin, Germany)*, August 2010.

[5] M.P. Cabasino, A. Giua, A. Paoli, and C. Seatzu. Decentralized diagnosis of Petri nets. In *Proc. 2010 American Control Conference (Baltimore, Maryland, US)*, June 2010.

[6] M.P. Cabasino, A. Giua, A. Paoli, and C. Seatzu. Decentralized diagnosability analysis of discrete event systems using Petri nets. In *Proc. of the IFAC 2011 World Congress (Milan, Italy)*, August 2011.

[7] M.P. Cabasino, A. Giua, M. Pocci, and C. Seatzu. Discrete event diagnosis using labeled Petri nets. An application to manufacturing systems. *Control Engineering Practice*, 19(9):989–1001, 2011.

[8] M.P. Cabasino, A. Giua, and C. Seatzu. Fault detection for discrete event systems using Petri nets with unobservable transitions. *Automatica*, 46(9):1531–1539, 2010.

[9] S.L. Chung. Diagnosing pn-based models with partial observable transitions. *International Journal of Computer Integrated Manufacturing*, 12 (2):158–169, 2005.

[10] O. Contant, S. Lafortune, and D. Teneketzis. Diagnosability of discrete event systems with modular structure. *Discrete Event Dynamic Systems*, 16(1):9–37, 2006.

[11] M.O. Cordier and L. Rozé. Diagnosing discrete-event systems: extending the "diagnoser approach" to deal with telecommunication networks. *Discrete Event Dynamic Systems*, 12(2):43–81, 2002.

[12] R. Debouk, S. Lafortune, and D. Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete-event systems. *Discrete Events Dynamic Systems*, 10(1):33–86, 2000.

[13] M. Dotoli, M.P. Fanti, A.M. Mangini, and W. Ukovich. On-line fault detection of discrete event systems by Petri nets and integer linear programming. *Automatica*, 45(11):2665–2672, 2009.

[14] H.E. Garcia and T.-S. Yoo. Model-based detection of routing events in discrete flow networks. *Automatica*, 41(4):583–594, 2004.

[15] S. Genc and S. Lafortune. Distributed diagnosis of place-bordered Petri nets. *IEEE Trans. Automation Science and Engineering*, 4(2):206–219, 2007.

[16] C. Hadjicostis. Probabilistic fault detection in finite-state machines based on state occupancy measurements. *IEEE Trans. Automatic Control*, 50(12):2078–2083, 2005.

[17] S. Jiang and R. Kumar. Diagnosis of repeated failures for discrete event systems with linear-time temporal logic specifications. *IEEE Trans. Automation Science and Engineering*, 3(1):47–59, 2006.

[18] G. Jiroveanu and R.K. Boel. A distributed approach for fault detection and diagnosis based on time Petri nets. *Mathematics and Computers in Simulation*, 70(5), 2006.

[19] J. Lunze. State observation and diagnosis of discrete-event systems described by stochastic automata. *Discrete Event Dynamic Systems*, 11(4):319–369, 2001.

[20] A. Paoli and S. Lafortune. Safe diagnosability for fault tolerant supervision of discrete event systems. *Automatica*, 41(8):1335–1347, 2005.

[21] Y. Pencole and M.O. Cordier. A decentralized model-based diagnostic tool for complex systems. *International Journal on Artificial Intelligence Tools*, 11(3):327–346, 2002.

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

[23] R. Su, W.M. Wonham, J. Kurien, and X. Koutsoukos. Distributed diagnosis for qualitative systems. In *Proc. 6th Work. on Discrete Event Systems (Zaragoza, Spain)*, 2002.

[24] Y. Wang, T.-S. Yoo, and S. Lafortune. Diagnosis of discrete event systems using decentralized architectures. *Discrete Event Dynamic Systems*, 17(2), 2007.