Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Distributed Diagnosis for Petri Nets models with unobservable interactions via common places

George Jiroveanu and René K. Boel

Abstract— **In this paper we consider the case of a large plant comprising different local sites. At each site, a local diagnoser must provide the diagnosis of the site based on the local plant model (a Petri Net), the local observation and the information exchanged with its neighbors. The communication between the local diagnosers is not event-driven and the interactions between the local sites (modeled as common places) are considered unobservable (tokens can enter and exit unobservably the local Petri Net models). For this general setting we present in this paper an algorithm that allows the local diagnosers to recover completely the results of a centralized diagnoser after the completion of a communication protocol.**

I. Introduction

Essential for all industrial activities the diagnosis of a plant answers the questions: *"Did a fault happen in the plant ?"* (fault detection) and *"Where did it happen ?" / "Which kind of fault happened if any ?"* (fault isolation) [12]. Additionally the diagnosis may be required for answering: *"How the fault happened ?"* (explanations [9]) and *"Which are the consequences of the fault occurrence ?"*.

The approach we follow in this paper for performing diagnosis is model based, with the plant model given as a Petri Net. Reasoning about the past evolution of the plant is based on observing a subset of observable events (whose occurrence is always reported). The faults that must be detected, are modeled as unobservable transitions. This approach assumes that the model and initial state are perfectly known and that observations are always correctly received.

In this paper we consider the case of a large plant consisting several interacting sub-systems (sites). Each sub-system is modeled as a Petri Net (PN). The interactions are represented by token passing via border (common) places from one sub-system to another [6], [13].

Since the centralized plant diagnosis [11] is usually not feasible for large plants, we assume that the knowledge about the model of the plant as well as the observation are distributed among a set of diagnoser agents (d-agents), with one d-agent located at each site. A d-agent knows the local sub-system (the PN model), receives the local site observation and it has the possibility to exchange limited information with its neighbors.

The general setting we consider is as follow. The communication between the d-agents is not event-driven i.e. the time the information exchange is allowed does not necessarily depend on the observations. It is also required that each d-agent performs a preliminary local diagnosis (PLD), based only on the local observations and the local model in absence of any external information.

In this distributed setting the main dif£culty in designing an algorithm to compute the plant diagnosis is the lack of knowledge on the marking of places where tokens can enter the local PN model. Basically a d-agent must estimate the behavior of the local site, the "marking" of the border places being uncertain [7].

The case of nondeterministic observation of the interactions was considered in [3] and [6]. Nondeterministic observation means that the input transitions and output transitions of any common place between any two subsystems are observable but there may be different observable transitions emitting the same label when they £re.

The case of unobservable interactions between the local sites (components) of the plant is studied in [2], [12] where the plant-model is given as a network of interacting automata. The proposed solution is to compute preliminary over-estimations for each local component that are checked for consistency by communication between the d-agents.

For PN models this solution is dif£cult to apply. It would require for each input place of a sub-system the computation of an upper bound on the number of tokens that can enter, implying a global analysis of the plant model. Even though this might be possible (by translating the PN model into a network of communicating automata), the approach based on upper bounds becomes utterly infeasible for PN models whenever the plant structure changes often (i.e. components are plugged in/out), this requiring the recalculation of the (new) upper bounds.

In [4] we proposed for general PN models and unobservable interactions via common places a distributed algorithm that allows the d-agents to detect the state F (*"a fault happened for sure"* [11]) whenever a centralized d-agent would detect the state F.

In this paper we adapt the algorithm presented in [4] s.t. the centralized diagnosis result is *completely recovered*: right after the communication protocol has been executed the d-agents also report the same faults that could have happened, as would the centralized diagnoser. The motivation is that the d-agents may be required from time to time to correlate perfectly their local estimates of the local states.

Without affecting the generality we consider here the case of only two interacting sites, deterministic labels for the observable events (while the unobservable events are silent),

This research was partially supported by IAP Program initiated by the Belgium State, Prime Minister's Of£ce for Science, Technology and Culture. G. Jiroveanu is supported by BOF - Ghent University.

The authors are with SYSTeMS Research Group, Ghent University, Technologiepark 914, Zwijnaarde 9052, Belgium {george.jiroveanu,rene.boel}@ugent.be

0-7803-9568-9/05/$20.00 ©2005 IEEE

and a global clock that governs the process (allowing for temporal ordering of the events observed in different sites).

The paper is organized as follow. Section II revises PNs notions and introduces the notation. In Section III we formally present the setting. In Section IV we present an algorithm for computing local preliminary diagnosis and in Section V we show how by communicating the agents accomplish the centralized diagnosis results. Finally in Section VI we conclude the paper with some £nal remarks.

II. Petri Nets

A Petri Net is a structure $N = \langle P,T,Pre,Post\rangle$ where P denotes the set of $\sharp P$ places, T denotes the set of $\sharp T$ transitions, and $Pre : P\times T \to \mathbb{N}$ and $Post : T\times P \to \mathbb{N}$ are the *pre-* and *post-incidence function* that specify the arcs. We use the standard notations: $p^\bullet$, ${}^\bullet p$ for the set of input, respectively output transitions of a place; similarly ${}^\bullet t$ and $t^\bullet$ denote the set of input places to t, and the set of output places of t respectively. A *marking* M of a PN is represented by a $\sharp P$-vector, $M : P \to \mathbb{N}$, that assigns to each place of N a non-negative number of tokens.

Denote by $\mathscr{L}_N(M_0)$ the set of all possible traces of PN $\langle N,M_0\rangle$, where a trace τ in $\langle N,M_0\rangle$ is de£ned as: $\tau = M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots \xrightarrow{t_k} M_k$, where for $i = 1\dots k$, $M_{i-1} \geq Pre(\cdot,t_i)$. $M \xrightarrow{\tau} M_k$ denotes that the enabled sequence τ may £re at M yielding M_k. Denote by $R_N(M_0)$ the set of all the reachable markings from M_0

The transition set T is partitioned into the disjoint subsets of observable T_o, and unobservable T_{uo} transitions. Then T_F denotes the set of faulty events whose occurrence must be detected ($T_F \subset T_{uo}$).

Denote by T^* the Kleene closure of the set T and by ε the empty string. Let $s \in \mathscr{L}_N(M_0) \subseteq T^*$. The projection $\Pi : \mathscr{L}_N(M_0) \to T_o^*$ (also denoted $\Pi_{T_o}(\mathscr{L}_N(M_0))$) is de£ned as: i) $\Pi(\varepsilon) = \varepsilon$; ii) $\Pi(t) = t$ if $t \in T_o$; iii) $\Pi(t) = \varepsilon$ if $t \in T_{uo}$; iv) $\Pi(st) = \Pi(s)\Pi(t)$ for $s \in \mathscr{L}_N(M_0)$ and $t \in T$.

For a set or a multiset X, 2^X is the set of all the sub-sets of X. Given $f : X \to Y$ and $A \subseteq X$ then $f(A) = \bigcup_{x\in A} f(x)$. Throughout the paper we treat a marking M as a vector or as a multi set of tokens.

De£nition 1: Given a PN N, $\wp = p_0t_1\dots t_np_n$ is a non-trivial unobservable elementary path in N if: i) $n > 0$; ii) $t_{q+1} \subseteq p_q^\bullet \cap {}^\bullet p_{q+1}$ for $q = 1,\dots,n$; iii) $t_q \in T_{uo}$ for $q = 1,\dots,n$.

De£nition 2: An unobservable elementary circuit (*uec*) denoted ζ is an unobservable elementary path $\wp$ that comprises different transitions and different places excepting the initial place p_0 and the £nal place p_n that are the same. If there is a place p_j $(0 \leq j \leq k)$ of an *uec* ζ that has more than one outgoing transition ($| p_j^\bullet | \geq 2$) we say that ζ is an *uec* with choice places and is denoted *uecwcp*.

An observed event is denoted t^o while a sequence of observed events is $O = t_1^o \dots t_\kappa^o$. The time an observable event t_q^o $(1 \leq q \leq \kappa)$ happened is denoted $\theta_{t_q^o}$. We assume that a global clock governs the overall process in the sense that $\theta_{t_q^o} < \theta_{t_w^o}$ when $q < w$.

III. The setting

We consider the distributed plant description as follow:

i) $N = N_1 \cup N_2$ where $N = \langle P,T,Pre,Post\rangle$ and for $i = 1,2$ $N_i = \langle P_i,T_i,Pre_i,Post_i\rangle$
ii) $P = P_1 \cup P_2$, $P_1 \cap P_2 = P_{12}$, $P_{12} \neq \emptyset$
iii) $T = T_1 \cup T_2$, $T_1 \cap T_2 = \emptyset$
iv) $Pre_i = Pre\,|_{N_i}$, $Post_i = Post\,|_{N_i}$ $i = 1,2$
v) $P_{12} = IN_i \cup OUT_i$, $IN_i \cap OUT_i = \emptyset$
vi) $IN_i = OUT_j = \left\{p \in P_{12} \mid p^\bullet \subseteq T_i \wedge {}^\bullet p \subseteq T_j\right\}$
vii) $IN_j = OUT_i = \left\{p \in P_{12} \mid {}^\bullet p \subseteq T_i \wedge p^\bullet \subseteq T_j\right\}$
viii) N is structurally bounded w.r.t. the unobservable evolution i.e. $\forall M \in \mathbb{N}^{\sharp P} \wedge \forall \sigma_{uo} \in T_{uo}^* : M \xrightarrow{\sigma_{uo}} M' \Rightarrow M' \ngtr M$

For simplicity we assume IN_i and OUT_i disjunct and $M_{0_{12}} = 0$ ($M_{0_{12}} = M_0(P_{12})$). When £red, an observable transition $t \in T_o$ emits a deterministic label $\delta(t)$ (i.e. $\delta(t_1) = \delta(t_2) \Rightarrow t_1 = t_2$), whereas an unobservable event does not emit anything ($\forall t \in T_{uo} \Rightarrow \delta(t) = \varepsilon$).

Given a marking $M_i \in \mathbb{N}^{\sharp P_i}$, denote by M_i, M_{IN_i} and M_{OUT_i} the marking of the places $P_i \setminus P_{ij}$, IN_i and OUT_i respectively.

At each site there is one d-agent Ag_i $(i = 1,2)$. The *a priori* knowledge of the local plant a d-agent Ag_i has (denoted K_{Ag_i}) comprises the local site model N_i, the set of events T_o^i whose occurrence it can observe (and are always observable) and the initial local state M_{0_i}; $K_{Ag_i} = \langle N_i,T_o^i,M_{0_i}\rangle$ (see Fig.1).

The information exchange between Ag_1 and Ag_2 is not event-driven. Denote by θ_c the £rst time the agents communicate. Denote by $O_{\theta_c}^i = t_{1_i} \dots t_{\kappa_i}$ the (possible empty) sequence of observed events recorded at the local site i by the time θ_c. The global observation O_{θ_c}, that a centralized agent AG ($K_{AG} = \langle N,T_o,M_0\rangle$) would receive by the time θ_c is $O_{\theta_c} = O_{\theta_c}^1 \otimes_{gc} O_{\theta_c}^2$ ($\otimes_{gc}$ states for the interleaving of the local observations $O_{\theta_c}^i$ according to a global clock).

The assumption made in this paper is that the process "stops" at time θ_c when the d-agents are required to achieve their goal by exchanging limited information in several consecutive messages. For $k = 0,\dots,K_{max}$, $Msg_{i\to j}^k$ denotes the k^{th} message sent by Ag_i to Ag_j at the time θ_c. Notice that there is no message before θ_c hence $Msg_{i\to j}^0 = \emptyset$. The local d-agents will not consider observations recorded after the time θ_c. A communication round with $k \geq 1$ implies that Ag_1 and Ag_2 exchange $Msg_{1\to 2}^k$, $Msg_{2\to 1}^k$ simultaneously. The consideration of asynchronous exchange of information brings nothing new but some more notation.

The set of (global) explanations of the received observation O_{θ_c} and the set of estimated states are de£ned as:

$$\begin{aligned} \mathscr{L}_N(O_{\theta_c}) &= \left\{\tau \in \mathscr{L}_N(M_0) \mid \Pi_{T_o}\tau = O_{\theta_c}\right\} \\ \mathscr{M}_N(O_{\theta_c}) &= \left\{M \mid \exists \tau \in \mathscr{L}_N(O_{\theta_c})\ s.t.\ M_0 \xrightarrow{\tau} M\right\} \end{aligned} \tag{1}$$

Eq.1 states that a global explanation τ is a possible evolution $\tau \in \mathscr{L}_N(M_0)$ that obeys the observation O_{θ_c}.

Consequently the centralized diagnosis at the time θ_c (denoted $D(\theta_c)$) results by projecting $\mathscr{L}_N(O_{\theta_c})$ onto the set of fault events T_F:

$$D(O_{\theta_c}) = \left\{\sigma_f \mid \sigma_f = \Pi_{T_F}\tau \wedge \tau \in \mathscr{L}_N(O_{\theta_c})\right\} \tag{2}$$

Given $D(O_{\theta_c})$ denote by $D_i(O_{\theta_c})$ $(i = 1,2)$ the centralized diagnosis for site i:

$$D_i(O_{\theta_c}) = \left\{\sigma_{f_i} \mid \sigma_{f_i} = \Pi_{T_i}\sigma_f \wedge \sigma_f \in D(O_{\theta_c})\right\} \qquad (3)$$

IV. Local preliminary diagnosis

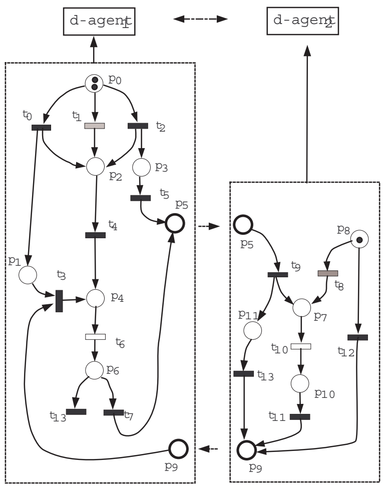

Fig. 1.

Consider the distributed architecture shown in Fig.1. The locally observable transitions are t_6 for N_1 on the left and t_{10} for N_2 on the right. The fault events whose occurrence should be detected are t_1 and t_8. Let the local observations at the time θ_c be $O^1_{\theta_c} = t_6t_6$ at site 1 and $O^2_{\theta_c} = t_{10}$ at site 2 while the global order is $O_{\theta_c} = t_6t_{10}t_6$.

Before starting to communicate Ag_2 does not know the marking of the input place p_5 and it must perform a local calculation before receiving external information. A solution would be to consider on the input place p_5 the maximum number of tokens that could have come from the beginning of the process up to the time θ_c but this may lead to local calculations of the same magnitude as the global calculation of the plant [13] that further increase the amount of information exchanged.

In [4] we have proposed for the same setting to calculate the minimum number of tokens required to minimally explain the local observation $O^i_{\theta_c}$. The computation is based on a backward reachability algorithm that starts from the observed event computing the minimal marking required on the input places IN_2 of N_2. It results that observing t_{10}, the minimal number of tokens required to have come via p_5 is $M(p_5) = 0$ if the local minimal explanation is $\tau^1_2 = t_8t_{10}$, or $M(p_5) = 1$ if $\tau^2_2 = t_9t_{10}$.

Here we extend this algorithm in the following way. Assume that Ag_2 has minimally explained the local observation as above. Based on the minimal explanations τ^1_2, τ^2_2 it also derived the estimated states: $M^1_2 = \{m(p_{10}) = 1\}$ and $M^2_2 = \{m(p_8) = 1, m(p_{10}) = 1, m(p_{11}) = 1\}$ respectively. What Ag_2 knows is that for τ^1_2 the token in p_{10} is available after the time $\theta_{t_{10}}$, when t_{10} was observed, and that for τ^2_2 a token coming at p_5 before $\theta_{t_{10}}$ is required and a tokens in p_{10} is available after $\theta_{t_{10}}$ and a token in p_{11} is available after the time t_9 has £red that is after the time the required tokens in p_5 has entered.

To take into account the timing information induced by the observation we associate with each token in a marking its rank as follows. Let a linear constraint be of the form $r \sim c$ with $\sim \in \{<;>\}$ where c is a constant ($ct \in \mathbb{R}^+$), a variable ($c \in \mathbb{R}^+$) or a conjunction of linear constrains.

Initially all the tokens in M_0 are considered as produced by the starting event t_{start} ($\theta_{t_{start}} = 0$) and they are given the rank $r(p) > 0$. Assume that at time θ_t a transition t £res and puts a token in a place p. This token is given the rank $\{r > \theta_t\}$. E.g. in M^2_2 $m(p_{10}) = \left\{r(p_{10}) > \theta_{t_{10}}\right\}$ where for simplicity we denote a token in a place by its timing constraint. Then any unobservable transition produces tokens having the rank $Max_{p_i \in {}^\bullet t}(c'_i) < r(p) < Max_{p_i \in {}^\bullet t}(c''_i)$ where $c'_i < r(p) < c''_i$ is the timing constraint of the tokens in the input places of t that are consumed by £ring t.

It means that depending on the tokens that are consumed £ring a transition, the newly produced tokens result in general with different ranks. Assume that the input places p_1 and p_2 of a transition t contain each two tokens as follow $m(p_1) = \left\{r(p_1) > \theta_{t^o_1}, r(p_1) > \theta_{t^o_3}\right\}$ and $m(p_2) = \left\{r(p_2) > \theta_{t^o_2}, r(p_2) > \theta_{t^o_4}\right\}$. Then t can £re in different ways e.g. consuming $r(p_1) > \theta_{t^o_1}$ and $r(p_2) > \theta_{t^o_4}$ and producing tokens having the rank $r(p) > \theta_{t^o_4}$ for $p \in t^\bullet$ (if $\theta_{t^o_4} > \theta_{t^o_1}$). Notice that we do not distinguish two tokens having the same linear constraint e.g. $r \sim_1 c_1$ and $r \sim_2 c_2$ where $\sim_1 = \sim_2$ and $c_1 = c_2$ (component wise). Moreover we say that a linear constraint is saturated if $\forall c_i$ s.t. $\wedge_i(r \sim c_i)$, $c_i = ct \in \mathbb{R}^+$ and not saturated otherwise.

The reason we set a linear constraint depending on variables is as follow. Consider again Ag_2 observing t_{10} as before. Thus $\tau^2_2 = t_9t_{10}$ requires a token in p_5 s.t. $\theta_{t_{10}} > r(p_5)$. Then $M^2_2 = \left\{r(p_8) > 0, r(p_{10}) > \theta_{t_{10}}, r(p_{11}) > r(p_5)\right\}$ that means that the token in p_{11} arrives latter than the token in p_5; the arrival time of a token in p_5 is a variable since there is no local knowledge Ag_2 has about its arrival.

Given the ranked marking M^2_2, Ag_2 estimates by forward search the ranked tokens that could have exited from the local site 2. E.g. τ^2_2 is extended by £ring either $\omega_1 = t_{13}$; $\omega_2 = t_8$; $\omega_3 = t_{11}$; $\omega_4 = t_{12}$; $\omega_5 = t_{13}t_8$; $\omega_6 = t_{13}t_{12}$; $\omega_7 = t_{13}t_{12}t_8$; .... Thus τ^2_2 extended by ω_7 will result in a ranked marking $M_{\omega_7} = \left\{r(p_9) < r(p_5), r(p_9) < \theta_{t_{10}}, r(p_7) > 0\right\}$.

A. Local preliminary calculation

In the following we present the backward computation of the set of minimal explanation $\mathscr{L}^{min}_N$ of the £rst observed event in the overall model N afterwards extending the approach to handle a sequence of observed events. Then we show how this method can be applied to a local model (e.g. N_i) whose marking is partially unknown (i.e. the marking of the input places IN_i is not know). Finally the set of minimal local explanations are extended for estimating the marking of the output places OUT_i.

De£ne $a \ominus b = a - b$ if $a \geq b$, and $a \ominus b = 0$ otherwise and extend "$\ominus$" to multisets in the natural manner [1].

De£nition 3: Backwards enabling rule: *A transition t is backward enabled in a marking $M \in \mathbb{N}^{\sharp P}$ iff $\exists p \in t^{\bullet}$ s.t. $M(p) \geq 1$.* Backwards £ring rule: *A backward enabled transition t in a marking $M \in \mathbb{N}^{\sharp P}$ £res backwards from M producing M' (denoted $M \overset{t}{\rightsquigarrow} M'$) where $M' = M \ominus Post(t,\cdot) + Pre(\cdot,t)$.*

A sequence of transitions $\tau = t_1 \ldots t_m$ is backward allowable from M (denoted $M \overset{\tau}{\rightsquigarrow} M'$) iff for $q = 1,\ldots,m$, $\tau_q = t_1 \ldots t_{q-1}$ and t_q is backward enabled in M'' where $M \overset{\tau_q}{\rightsquigarrow} M''$.

De£nition 4: Given a PN N, consider a marking $M \in \mathbb{N}^{\sharp P}$. Then M is covered by M' iff $\exists \sigma \in \mathscr{L}_N(M')$, and $M' \xrightarrow{\sigma} M''$ $\wedge$ $M'' \geq M$.

Proposition 1: Given a PN $\langle N, M_0 \rangle$ and a marking M, then M is covered by M_0 iff $\exists M' \leq M_0$ s.t $M \overset{\sigma}{\rightsquigarrow} M'$.

Denote by $BC_N(M)$ the set of all the markings that cover M: $BC_N(M) = \left\{ M' \mid M \overset{\sigma}{\rightsquigarrow} M' \right\}$. Then we have that M is *covered unobservably by M_0* iff $\exists M' \in BC_N(M)$ s.t $M' \leq M_0$ and $M \overset{\sigma_{uo}}{\rightsquigarrow} M'$ where $\sigma_{uo} \in T_{uo}^*$. We denote by $BC_N^{uo}(M)$ the set of markings that unobservablly cover the marking M.

Let a PN $\langle N, M_0 \rangle$ and the £rst observed transition t_1^o. Denote $M_{t_1^o} = Pre(\cdot, t_1^o)$. Then we have:

$$\begin{aligned} &\mathscr{L}_N^{min}(t_1^o) = \left\{ \tau = \sigma_{uo} t_1^o \mid M_{t_1^o} \overset{\sigma_{uo}}{\rightsquigarrow} M' \wedge M' \in BC_N^{uo}(M_{t_1^o}) \cap 2^{M_0} \right\} \\ &\mathscr{M}_N^{min}(t_1^o) = \left\{ M_\tau \mid M_0 \xrightarrow{\tau} M_\tau \wedge \tau \in \mathscr{L}^{min}(t_1^o) \right\} \end{aligned} \tag{4}$$

We have that $\mathscr{M}_N^{min}(t_1^o) \subseteq \mathscr{M}_N(t_1^o)$ and:

$$\bigcup_{M \in \mathscr{M}_N^{min}(t_1^o)} R_N(M) \equiv \bigcup_{M' \in \mathscr{M}_N(t_1^o)} R_N(M') \tag{5}$$

$\mathscr{L}_N^{min}(t_1^o)$ results after intersecting $BC_N^{uo}(M_{t_1^o})$ with 2^{M_0}. Our problem looks similar to the model-checking problem, where one checks whether a marking M is covered/reachable from an initial marking M_0 [1], [5]. The difference is that we should compute the set of explanations ($\mathscr{L}^{min}$) that assures completeness regarding the future behavior (see Eq.5) whereas in model-checking it suf£ces to compute the existence of a solution (one explanation).

Problem: check whether M is covered by M_0 or not. The standard termination condition for searching backward along a leaf is: *"If $M \overset{\sigma'}{\rightsquigarrow} M' \wedge M' \overset{\sigma''}{\rightsquigarrow} M''$ and $M'' \geq M'$ then* abort the search along the leaf.*"*

We cannot use this termination condition since its application may result in omitting the calculation of some explanations in $\mathscr{L}_N^{min}$. We illustrate this by the following example.

Example 1: Consider the PN shown in Fig. 2 where the dotted lines emerging from t_4 and p_4 indicate that the displayed part is a sub-net of a large PN model. When the event t_4 (the only observable event in the displayed part of the net) was observed, then $M_{t_4} = \{m(p_3)\}$ and we can derive $\tau_1 \in \mathscr{L}_N^{min}(t_4)$, $\tau_1 = t_5 t_0 t_1 t_4$. Notice that there is another minimal explanation $\tau_2 = t_5 t_0 t_1 t_2 t_3 t_5 t_0 t_1 t_4$ that would not have been found by applying TC.

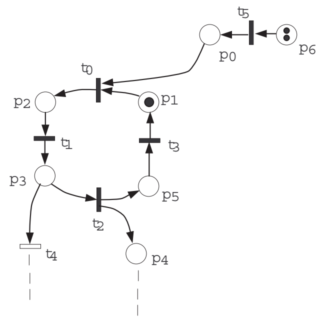

Fig. 2.

The following two results allow us to discard backward traces that will not result in minimal traces.

Proposition 2: Given $\langle \mathscr{N}, M_0 \rangle$ and a marking M that is not covered by M_0 then $\forall M' > M$, M' is not covered by M_0.

Proposition 3: Given $\langle \mathscr{N}, M_0 \rangle$ and a marking M then:

$$BC_N^{uo}(M) \cap 2^{M_0} \neq \emptyset \text{ if } \forall M', M' < M, BC_N^{uo}(M') \cap 2^{M_0} \neq \emptyset \tag{6}$$

Thus, during computation, a set of markings to be processed (say SET) is maintained, where at the current step a marking M is processed if there is not any other marking M' in SET s.t. $M' < M$ (Proposition 2). The markings M' that are not minimal elements in SET w.r.t. $<$ are "stopped" until one has checked that $\forall M$, $M \in SET \wedge M < M' \Rightarrow \exists \sigma_{uo}$ s.t. $M \overset{\sigma_{uo}}{\rightsquigarrow} M'' \wedge M'' \leq M_0$ (Proposition 3).

Then the method extends to a sequence of observed events $O = t_1^o t_2^o$ in a straightforward manner. E.g. for the second observed event t_2^o the backward algorithm applies for all $M_1^{min} \in \mathscr{M}_N^{min}(t_1^o)$ (see Eq.4) instead of M_0. The minimal explanations for $\tau^o = t_1^o t_2^o$ results by concatenating $\tau_1 \tau_2$ where $\tau_1 \in \mathscr{L}_N^{min}(t_1^o)$, $M_0 \xrightarrow{\tau_1} M_1^{min}$ and $\tau_2 \in \mathscr{L}_N^{min}(t_2^o)$, $M_1^{min} \xrightarrow{\tau_2} M_2^{min}$. Notice that the completeness is guaranteed by Eq. 5

Consider in the following the backward calculation for the distributed setting (e.g. Ag_i and K_{Ag_i}) where the marking of some places (e.g. IN_i) is unknown.

Example 2: Consider again the PN in Fig. 2 where t_4 was observed but this time consider that p_0 is an input place on the border. In this case the computation does not terminate since running in£nitely backward the uecwcp $\zeta := t_0 t_1 t_2 t_3$ there will be a minimal local explanation requiring an in£nite number of tokens entering at the input place p_0 before t_4 was observed.

This motivates the following structural assumption we must impose in order to assure that the local backward search terminates without assuming in the input places in£nite number of (ranked) tokens.

Given an *uecwcp* ζ, denote by Υ_ζ the set of limiting places of ζ: $\Upsilon_\zeta \overset{\triangle}{=} \{p \mid p \notin \zeta \wedge \exists t \in \zeta \ s.t.\ p \in {}^{\bullet}t\}$. A place $p \in \Upsilon_\zeta$ is a limiting places of ζ since every complete execution of ζ consumes one token from p. For $M_{\Upsilon_\zeta} \neq \emptyset$ let $M_{\Upsilon_\zeta} = \left\{ m(p) = 1 \mid p \in \Upsilon_\zeta \right\}$.

Assumption 1: For any local model N_i $(i = 1, 2)$ and for any uecwcp *ζ_i, there does not exist an executable sequence* *of unobservable transitions σ_{uo_i} with initial marking the marking M that has tokens only in the input places IN_i ($M(p) = 0$ for $p \notin IN_i$) s.t. by £ring from M, σ_{uo_i} produces a marking M' greater than the limiting marking of ζ_i, $M_{\Upsilon_{\zeta_i}}$.*

$\nexists \sigma_{uo} \in T^*_{uo_i}$ s.t. ($M_{\Upsilon_{\zeta_i}} \overset{\sigma_{uo_i}}{\rightsquigarrow} M$) $\wedge$ ($M(p) \neq 0 \Rightarrow p \in IN_i$).

Then it is easy to see that if the PN model $\langle N, M_0\rangle$ is bounded w.r.t. the unobservable evolution (item *viii* in setting) and Assumption 1 is satis£ed the backward calculation terminates after £nitely many steps.

Thus for any observed sequence $O^i_{\theta_c} = t^o_{1_i} \dots t^o_{\kappa_i}$ a local agent Ag_i derives backward the set of local minimal explanations $\mathscr{L}^{min}_{N_i}(O^i_{\theta_c})$, the set of minimum required tokens $\mathscr{M}_{IN_i}$ and the set of estimated states $\mathscr{M}^{min}_{N_i}$ where:

$$\begin{aligned}\forall \tau_i \in \mathscr{L}^{min}_{N_i}(O^i_{\theta_c}) \Rightarrow \Pi_{T_{o_i}} \tau_i = O^i_{\theta_c} \wedge \exists M_{IN_i} \in \mathscr{M}_{IN_1} \wedge \\ \exists M^{min}_i \in \mathscr{M}^{min}_{N_i} \;\; s.t. \;\; M_{0_i} \uplus M_{IN_i} \xrightarrow{\tau_i} M^{min}_i\end{aligned} \tag{7}$$

Notice that $M_{IN_i} \in \mathscr{M}_{IN_i}$ and $M^{min}_{N_i} \in \mathscr{M}^{min}_{N_i}$ are multi-sets of *ranked tokens* where each token has a rank (a timing constraint) associated with it.

B. Information exchange and local updates

Consider bellow that for the observed sequence $O^i_{\theta_c}$ received locally by the time θ_c Ag_i has derived: $\mathscr{L}^{min}_{N_i}(O^i_{\theta_c})$ $\mathscr{M}^{min}_{IN_i}(O^i_{\theta_c})$ $\mathscr{M}^{min}_{N_i}(O^i_{\theta_c})$. In order to simplify the notation we drop the indexes that are obvious from the context e.g. $\mathscr{L}^{min}_{N_i}(O^i_{\theta_c})$ will be denoted in short $\mathscr{L}^{min}_i$.

Since at θ_c Ag_i will be communicating with its neighbor, it should also calculate just prior to θ_c an estimate of the marking of the output places OUT_i to be included in the £rst message that is sent. To do this Ag_i computes the "forward unobservable extensions" of the minimal explanations computed as $\mathscr{L}^{min}_i$ deriving $\mathscr{L}^{ext}_i$ by a forward search starting from every ranked marking $M^{min}_i \in \mathscr{M}^{min}_i$.

$$\begin{aligned}\mathscr{L}^{ext}_i(O_i) \overset{\triangle}{=} \Big\{ \tau_i \omega_i \mid \exists M_{IN_i} \in \mathscr{M}^{min}_{IN_i} \wedge \\ \wedge M_{0_i} \uplus M_{IN_i} \xrightarrow{\tau_i} M^{min}_i \wedge \omega_i \in \mathscr{L}_{N_i}(M^{min}_i) \wedge \omega_i \in T^*_{uo} \Big\}\end{aligned} \tag{8}$$

where ω_i represents an unobservable possible continuation from an estimated state M^{min}_i reached by considering a minimal explanation $\tau_i \in \mathscr{L}^{min}_i$ of the observed event $O^i_{\theta_c}$ providing the minimum number of tokens M_{IN_i} was received at the input placed IN_i. Hence we obtain:

$$\begin{aligned}\mathscr{M}^{ext}_i(O_i) \overset{\triangle}{=} \Big\{ M^{ext}_i \mid M_{0_i} \uplus M_{IN_i} \xrightarrow{\tau_i \omega_i} M^{ext}_i \wedge \\ \tau_i \omega_i \in \mathscr{L}^{ext}_i(O_i) \Big\}\end{aligned} \tag{9}$$

If we project the ranked markings $M^{ext}_i \in \mathscr{M}^{ext}_i$ on to the output places OUT_i we obtain the set: $\mathscr{M}^{ext}_{OUT_i} = \{M_{OUT_i} \mid M^{ext}_i \in \mathscr{M}^{ext}_i \wedge M_{OUT_i} = M^{ext}_i(OUT_i)\}$.

It results that having received the set of inputs $\mathscr{M}_{IN_i}$ the local observation is explained by $\mathscr{L}^{ext}_i(O^i_{\theta_c})$ that has as result the set of "output" markings: $\mathscr{M}_{OUT_i}$.

The £rst message sent by Ag_i to Ag_j is $Msg^1_{i \to j} = \langle \mathscr{M}_{IN_i}, \mathscr{M}_{OUT_i}, \phi^1_i \rangle$ where $\phi^1 \subseteq \mathscr{M}_{IN_i} \times \mathscr{M}_{OUT_i}$ is the input-output correlation relation (non-empty only for compatible combinations of markings at the input-output places). Notice that when Ag_i sends Ag_j the message $Msg^1_{i \to j}$ at the same time Ag_j sends to Ag_i the £rst message $Msg^1_{j \to i}$ derived in a similarly way.

De£nition 5: Given two local traces $\tau_i \in T^*_i$, $\tau_j \in T^*_j$, (τ_i, τ_j) is a consistent pair of local traces if $\exists \tau \in \mathscr{L}_N(M_0)$ s.t. $\Pi_{T_i} \tau = \tau_i$ and $\Pi_{T_j} \tau = \tau_j$.

We "recover" the entire set of pairs of consistent local traces $\mathscr{L}_1(O_c) \times \mathscr{L}_2(O_c)$ ($\mathscr{L}_i(O_c) = \Pi_{T_i}(\mathscr{L}_N(O_c))$, $i = 1, 2$) in the following way.

For each pair $\tau_i \in \mathscr{L}^{ext}_i$ and $\tau_j \in \mathscr{L}^{ext}_j$ we check if (τ_i, τ_j) is consistent and then if either τ_i or τ_j, or both can be extended for generating new pairs. Notice that in general $\mathscr{L}^{ext}_i(O^i_c) \nsubseteq \mathscr{L}_i(O_c)$ and $\mathscr{L}_i(O_c) \nsubseteq \mathscr{L}^{ext}_i(O^i_c)$. Then for a PN N structurally bounded w.r.t unobservable evolution (see *viii* in setting), the computation achieves a £x-point (no new pairs can be generated). All the pairs that are not consistent when the £x-point is achieved are dropped.

We illustrate the method by considering an arbitrary pair (τ_i, τ_j) where $M_{0_i} \uplus M_{IN_i} \xrightarrow{\tau_i} M_i \uplus M_{OUT_i}$ and $M_{0_j} \uplus M_{IN_j} \xrightarrow{\tau_j} M_j \uplus M_{OUT_j}$.

Let M_{OUT_j} be partitioned into the set of saturated ranked tokens $M^s_{OUT_j}$ and the set of unsaturated ranked tokens $M^{us}_{OUT_j}$ ($M_{OUT_j} = M^s_{OUT_j} \cup M^{us}_{OUT_j}$). Checking consistency implies to £nd an assignment between the ranked tokens required and ranked tokens produced such that all the ranked tokens required (M_{OUT_i}, M_{OUT_j}) are substituted by produced ranked tokens (M_{IN_i}, M_{IN_j}) and all the timing constraints are saturated (satis£ed).

Def£nition 6: De£ne the assignment function $\psi_i : M_{IN_i} \to M^s_{OUT_j} \cup \{\varepsilon\}$ where for $m_i \in M_{IN_i}$ we have:

$$\psi_i(m_i) = \begin{cases} m_j \in M^s_{OUT_j} \;\; \text{if}\; m_j := r(p) > c_j \wedge c_j < c_i \\ or \;\; \varepsilon \end{cases} \tag{10}$$

and for $m_{i_1} \neq m_{i_2}$, $\psi_i(m_{i_1}) = \psi_i(m_{i_2}) \Rightarrow \psi(m_{i_1}) = \varepsilon$. Similarly de£ne $\psi_j : M_{IN_j} \to M^s_{OUT_i} \cup \{\varepsilon\}$. Then denote by $\psi = (\psi_i, \psi_j)$ the assignment function of the common place marking and by Ψ the entire set of assignment functions.

Notice that because of the unobservable interactions there may be circular dependencies between the ranked tokens s.t. substituting e.g. a required ranked token at $r(p) \in M_{IN_i}$ with a saturated ranked token $r'(p) \in M_{OUT_j}$ may lead to further saturation of ranked tokens at $r''(p) \in M_{OUT_i}$. Moreover for assuring the completeness we must also consider the case when there is no substitution for the required tokens but $\psi(M_{IN_i}) = \varepsilon$ although there were enough ranked tokens for substituting the required tokens. All the saturated tokens that are not assigned to required tokens are considered as new entered tokens $M^{new}_{IN_i}(\psi)$ under the assignment ψ. Notice that whenever $M^{new}_{IN_i} \neq \emptyset$ then τ_i can be extended £ring new transitions that become possible only after the arrival of the new ranked tokens $M^{new}_{IN_i}$.

Hence M_{IN_i} under the assignment $\psi \in \Psi$ is:

$$M_{IN_i}(\psi) = M^a_{IN_i}(\psi) \cup M^{ua}_{IN_i}(\psi) \cup M^{new}_{IN_i}(\psi) \tag{11}$$

$M^a_{IN_i}(\psi) = \{m_i(p) \wedge m_j(p) \mid \psi_i(m_i) = m_j\}$ is the set of assigned tokens (notice that by substituting $m_i(p)$ with $m_j(p)$ we obtain ranked token $c_j < r(p) < c_i$). $M^{ua}_{IN_i}(\psi) =$

$\{m_i(p) : r(p) < c_i \mid \psi_i(m_i) = \varepsilon\}$ is the set of unassigned tokens and $M^{new}_{IN_i}(\psi) = \left\{m_j \in M^s_{OUT_j} \setminus \psi(M_{IN_i})\right\}$ is the set of new entered tokens.

Given two local explanations τ_i, τ_j and an assignment ψ we have that:

- the pair $(\tau_i, \tau_j)_\psi$ is consistent under the assignment ψ iff $M^{ua}_{IN_i}(\psi) = \emptyset$ and $M^{ua}_{IN_j}(\psi) = \emptyset$
- if either $M^{ua}_{IN_i}(\psi) \neq \emptyset$ or $M^{ua}_{IN_j}(\psi) \neq \emptyset$ then the pair $(\tau_i, \tau_j)_\psi$ is inconsistent
- if either $M^{new}_{IN_i}(\psi) \neq \emptyset$ or $M^{new}_{IN_j}(\psi) \neq \emptyset$ then the pair $(\tau_i, \tau_j)_\psi$ is extendable

For an extendable pair $(\tau_i, \tau_j)_\psi$ denote $\Delta(\tau_i^\psi) = \left\{\tau_i \omega_i \mid \omega_i \in \mathscr{L}_i(M_{\tau_i} \uplus M^{new}_{IN_i}) \cap T^*_{uo_i}\right\}$ that is the set of all the traces that extend the extendable trace τ_i^ψ by £ring unobservable strings ω_i with the new entered tokens $M^{new}_{IN_i}$.

Then denote $\Delta\mathscr{L}_i = \left\{\Delta(\tau_i^\psi) \mid \forall(\tau_i, \tau_j) \wedge \forall \psi \in \Psi\right\}$ the local update of Ag_i after receiving the £rst message.

For $\Delta\mathscr{L}_i \neq \emptyset$ Ag_i calculates $\Delta\mathscr{M}_{OUT_i}$ and then if $\Delta\mathscr{M}_{OUT} \neq \emptyset$ it sends the next message $Msg^2_{i\to j} = \langle \Delta\mathscr{M}_{IN_i}, \Delta\mathscr{M}_{OUT_i}, \psi^2_i \rangle$ otherwise Ag_i sends the termination message (it holds for two agents since with no more new tokens entered a d-agent cannot generate new more tokens). Notice that $\Delta\mathscr{M}_{IN_i} = \left\{M_{IN_i}(\psi) \mid \forall \tau_i^\psi \wedge \forall \psi \in \Psi\right\}$.

V. The algorithm

DD_Algo(Ag_i considered)
input: $\langle N_i, M_{0_i} \rangle$, $O^i_{\theta_c}$
output: $PLD_i(O^i_{\theta_c})$

```
k_i = 1, L_i^1 = L_i^ext, M_i^k = M_i^ext, M_{IN_i}^k, M_{OUT_i}^k
do
   send Msg_{i→j} = ⟨M_{IN_i}^k, M_{OUT_i}^k, φ_i^k⟩
   receive Msg_{j→i} = ⟨M_{IN_j}^k, M_{OUT_j}^k, φ_j^k⟩
   for each (τ_i, τ_j) ∈ L_i^k × L_j^k do
     for each ψ ∈ Ψ do
       if (τ_i, τ_j)^ψ consistent then L_i^cons = L_i^cons ∪ {τ_i}
       if (τ_i, τ_j)^ψ is N_i extendable then
       compute Δ(τ_i^ψ) and ΔL_i^k = ΔL_i^k ∪ Δ(τ_i^ψ)
     od
   od
   calculate ΔM_{IN_i}^k, ΔM_{OUT_i}^k
   if ΔM_{OUT_i}^k = ∅ then Msg_{i→j} = stop
   else k = k+1; M_{IN_i}^k = ΔM_{IN_i}^{k-1}; M_{OUT_i}^k = ΔM_{OUT_i}^{k-1}
while Msg_{j→i} = stop or Msg_{i→j} = stop
PLD_i(O_{θ_c}^i) = Π_{T_{F_i}}(L_i^cons)
```

Now for $p_0 \in P_{ij}$ denote by $c^i_\wp$ and $c^j_\wp$ how many times an unobservable oriented path $\wp$ crosses $P_i \setminus P_{ij}$ and $P_j \setminus P_{ij}$ respectively. Let $c_\wp = max(c^i_\wp, c^j_\wp)$ and denote $K_c = Max_{\wp \in N}(c_\wp)$. If $\exists \wp \in N$ s.t. $\wp$ is an *uec* then $K_c = \infty$ otherwise K_c is £nite.

Theorem 1: Consider a distributed description of the plant and an arbitrary distributed observation $O_{\theta_c} = O^1_{\theta_c} \otimes O^2_{\theta_c}$. Then the local preliminary diagnosis PLD^k_i $(i = 1,2)$ derived by the d-agent Ag_i at the time θ_c after the k^{th} communication round by running the algorithm DD_Algo *is such that:*

i) *if K_c is £nite then for $k \geq K_c$, $PLD^k_i(O^i_{\theta_c}) = D_i(O_{\theta_c})$*

ii) *if K_c is in£nite then $\exists$ $k_{max} \in \mathbb{N}^+$ £nite s.t. for $k \geq k_{\max}$, $PLD^k_i(O^i_{\theta_c}) = D_i(O_{\theta_c})$ where k_{max} depends on both the PN topology and the initial marking M_0.*

Proof: i), and ii) have a similar proof. First we prove that the computation achieves a £x-point. The proof relies on the structurally boundness see viii) in setting and is omitted because lack of space. Then the completeness is guaranteed by the way ψ is de£ned. $\square$

Example 3: Let the pair (τ^1_1, τ^2_1) with $\tau^1_1 = t_0 t_3 t_6 t_7 t_4 t_6$ and $\tau^1_2 = t_9 t_{10} t_{12}$. We have that $(\tau^1_1, \tau^1_2)^\psi$ are consistent but not extendable where $M^1_{IN_1} = \left\{r(p_9) < \theta_{t^1_6}\right\}$, $M^1_{OUT_1} = \left\{r(p_5) > \theta_{t^1_6}\right\}$; $M^1_{IN_2} = \left\{r(p_5) < \theta_{t_{10}}\right\}$, $M^1_{OUT_2} = \left\{r(p_9) > 0\right\}$ and $m_1 : r(p_9) < \theta_{t^1_6} \xrightarrow{\psi_1} m_2 : r(p_9) > 0$ and $m_2 : r(p_5) < \theta_{t_{10}} \xrightarrow{\psi_2} m_1 : r(p_5) > \theta_{t^1_6}$. The pair (τ^2_1, τ^2_2) with $\tau^2_1 = t_0 t_3 t_6 t_0 t_3 t_6$ and $\tau^2_2 = t_8 t_{10} t_{11}$ are not consistent since $M^2_{IN_1} = \left\{r(p_9) < \theta_{t^1_6}, r(p_9) < \theta_{t^2_6}\right\}$, $M^2_{OUT_1} = \emptyset$. $M^2_{IN_2} = \emptyset$, $M^2_{OUT_2} = \left\{r(p_9) > \theta_{t_{10}}\right\}$.

VI. Conclusions

This research is motivated by our interest in designing distributed algorithms for large plants where (e.g. because of sensors failure) unobservable inputs are sent/received between components placed in different sites. We have shown that by backward/forward search including linear timing constraints the centralized diagnosis result can be recovered. For increasing the ef£ciency of the local calculations one can use reachability methods based on unfoldings (backward [1] and forward unfolding [3],[10]). Further directions are the extension of this method (for reasonable models where $K_c \leq 2$) to time PN models and probabilistic analyze.

References

[1] P.A. Abdulla, S.P. Iyer and A. Nylen On Unfolding unbounded Petri Nets *Proceedings of Computer Aid Veri£cation 2000*, LNCS, vol.1855

[2] P.Baroni and G.Lamperti Diagnosis of large active systems. *Arti£cial Intelligence*, 110(2), 1999

[3] A. Benvensite, S. Haar et al. Distributed monitoring of concurrent and asynchronous systems *ATPN'03*, Eindhoven, 2003

[4] R. K. Boel and G. Jiroveanu Distributed Contextual Diagnosis for very large systems *WODES'04*, Reims, 2004

[5] A. Finkel, J-F. Raskin, et al. Monotonic extensions of Petri Nets: forward and backward search revisited *Technology Transfer*, 2001

[6] S. Genc and S. Lafortune Distributed diagnosis for DES using Petri Nets *ATPN'03*, Eindhoven, 2003

[7] A. Giua and C. Seatzu Observability of Place/Transition Nets *IEEE Transaction on Automatic Control*, 47(9), 2002

[8] G. Jiroveanu, R.K. Boel and B. Bordbar Contextual analysis of partially observable large Petri Net models *submitted to JDEDS*, 2004

[9] S. McIlraith Explanatory diagnosis: Conjecturing actions to explain observation *Conf. on Principles of Knowledge Representation*, 1998

[10] K. L. McMillan Using unfoldings to avoid the state space explosion problem in veri£cation of asynchronous circuits LNCS, vol. 663, 1992

[11] M. Sampath, R. Sengupa et.al Diagnosability of Discrete Event Systems *Report*, The University of Michigan, 1994

[12] R. Su Distributed diagnosis for Discrete Event Systems *PhD These*, University of Toronto, 2004

[13] A. Valmari Compositional analysis with place-bordered subnets *LNCS 815*, pp. 531-547, 1994