

Technical report 06-018

On-line diagnosis for time Petri nets*

G. Jiroveanu, B. De Schutter, and R.K. Boel

If you want to cite this report, please use the following reference instead:

G. Jiroveanu, B. De Schutter, and R.K. Boel, "On-line diagnosis for time Petri nets," *Proceedings of the 17th International Workshop on Principles of Diagnosis (DX-06)*, Burgos, Spain, 8 pp., June 2006.

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.51.19 (secretary)
fax: +31-15-278.66.79
URL: <http://www.dcsc.tudelft.nl>

*This report can also be downloaded via http://pub.deschutter.info/abs/06_018.html

On-line diagnosis for Time Petri Nets

G. Jiroveanu^{†,*} B. De Schutter[‡] R.K. Boel[†]

[†] EESA - SYSTeMS, University of Ghent, Belgium
{george.jiroveanu,rene.boel}@ugent.be

[‡] DCSC, Delft University of Technology, The Netherlands
b.deschutter@dcsc.tudelft.nl

Abstract

We derive in this paper on-line algorithms for fault diagnosis of Time Petri Net (TPN) models. The plant observation is given by a subset of transitions while the faults are represented by unobservable transitions. The model-based diagnosis uses the TPN model to derive the legal traces that obey the received observation and then checks whether or not fault events occurred. To avoid the consideration of all the interleavings of the concurrent transitions, the plant analysis is based on partial orders (unfoldings). The legal plant behavior is obtained as a set of configurations. The set of legal traces in the TPN is obtained solving a system of $(max, +)$ -linear inequalities called the characteristic system of a configuration. We present two methods to derive the entire set of solutions of a characteristic system, one based on Extended Linear Complementarity Problem and the second one based on constraint propagation that exploits the partial order relation between the events in the configuration.

1 Introduction

This paper deals with the diagnosis of TPNs. TPNs are extensions of untimed Petri Nets (PNs) where information about the execution delay of some operations is available in the model. In a TPN a transition can be fired within a given time interval after its enabling and its execution takes no time to complete. A trace in the plant comprises the transitions that are executed in the TPN model (the untimed support) as well as the time of their occurrence.

Since a transition can be executed at any time within an interval after it has become enabled, the state space of TPNs is in general infinite. Methods based on grouping states under a certain equivalence relation onto so called state classes were proposed in [2]. The state class graph was proved to be finite iff the net is bounded, thus infinite state spaces can be finitely represented and the analysis of TPN models is computable.

*Supported by a European Union Marie Curie Fellowship during his stay at Delft University of Technology (HPMT-CT-2001-0028). Currently with TRANSELECTRICA SA, Craiova, Romania.
e-mail: george.jiroveanu@transelectrica.ro

We consider the plant observation given by a subset of transitions whose occurrence is always reported. Moreover the time when an observed transition is executed is measured and reported according to a global clock. The unobservable events are silent, i.e. the execution of an unobservable transition is not acknowledged to the monitoring system. The faults are modeled by a subset of unobservable transitions.

The model-based diagnosis for TPNs comprises two stages. First the set of traces that are legal and that obey the received observation is derived and then the diagnosis result of the plant is obtained checking whether some or all of the legal traces include fault transitions.

The diagnosis of a TPN can be derived based on the computation of the state class graph as proposed in [5]. However the analysis of TPNs is not tractable even for models of reasonable size because of the interleaving of (unobservable) concurrent transitions.

Partial orders were shown to be an efficient method to cope with the state space explosion of untimed PNs because the interleaving of concurrent transitions is filtered out [4],[8]. They were also found applicable for the analysis of PN models where the time is considered as quantifiable and continuous parameter [1],[3].

In this paper we extend the results presented in [6],[7] presenting on-line algorithms for the diagnosis of TPNs based on partial orders. The plant analysis is based on time configurations (time-processes in [1]). A time configuration is an untimed configuration (a configuration in the net-unfolding of the untimed PN support of the TPN model) with a valuation of the execution times for its events. A time configuration is legal if there is a time trace in the original TPN that can be obtained from a linearization of the events of the configuration where the occurrence times of the transitions in the trace are identical with the valuation of their images in the time configuration. A linearization of the events in a configuration is a trace that comprises all the events of the configuration executed only once such that the partial order between the events in the configuration is preserved in the order in which they appear in the trace.

The on-line diagnosis algorithm that we propose works as follows. When the process starts we derive time traces in the TPN model up to the first discarding time. A discarding time is the time when in absence of any observation one can discard untimed support traces and it corresponds with

the smallest value of the latest time when an observable event could be forced to happen. The occurrence of an observable transition before the first discarding time is taken in to account eliminating traces that are not consistent with the received observation. Then the plant behavior is derived up to a next discarding time.

The set of all legal time traces in the original TPN can be obtained computing for each configuration the entire solution set of a system of $(\max, +)$ -linear inequalities called the characteristic system of the configuration.

The calculations involve time interval configurations. A time interval configuration is an untimed configuration endowed with time intervals for the execution of the events within the configuration. A time interval configuration is legal if for every event and for every execution time of the event within its execution time interval there exists a legal time configuration that considers the event executed at that time.

Thus, we need to derive for each configuration the entire solution set of its characteristic system. The naive approach to enumerate all the possible max-elements would imply to interleave concurrent events which is exactly what we wanted to avoid by using partial orders to represent the plant behavior. To cope with this difficulty we present two methods that avoid the explicit consideration of all the cases for each max-term in the characteristic system.

The first method uses the Extended Linear Complementarity Problem (ELCP) [10] for deriving the set of all solutions of the characteristic system of the configuration. The solution set can be represented as a union of faces of a polyhedron that satisfy a cross-complementarity condition.

The second method is based on constraint propagation and exploits the partial order relation between the events within the configuration. We derive for each untimed configuration a set of hyperboxes of dimension equal with the number of events within the configuration such that the union of all the subsets of solutions that are circumscribed by the hyperboxes is a cover of the solution set.

The paper is organized as follows. In Section 2 we provide definitions and the notation used in the paper. In Section 3 we formalize the diagnosis problem for TPNs models. The analysis of TPNs based on partial orders is described in Section 4. Section 5 and Section 6 present the two methods to derive the solution set of a characteristic system of a configuration and then in Section 7 we present the on-line diagnosis algorithm that we propose. The paper is concluded in Section 8 with final remarks and future work.

2 Notation and definitions

2.1 Petri nets

A Petri Net is a structure $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F)$ where \mathcal{P} denotes the set of $|\mathcal{P}|$ places, \mathcal{T} denotes the set of $|\mathcal{T}|$ transitions, and $F = Pre \cup Post$ is the incidence function where $Pre(p, t) : \mathcal{P} \times \mathcal{T} \rightarrow \{0, 1\}$ and $Post(t, p) : \mathcal{T} \times \mathcal{P} \rightarrow \{0, 1\}$ are the pre- and post-incidence function that specify the arcs.

We use the standard notations: \bullet^{\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

$|\mathcal{P}|$ -vector, $M : \mathcal{P} \rightarrow \mathbb{N}$, that assigns to each place of \mathcal{N} a non-negative number of tokens.

The set $\mathcal{L}_{\mathcal{N}}(M_0)$ of all legal traces of a PN, $\langle \mathcal{N}, M_0 \rangle$, with initial marking M_0 is defined as follows. A transition t is *enabled* at the marking M if $M \geq Pre(\cdot, t)$. Firing, an enabled transition t consumes $Pre(p, t)$ tokens in the input places $p \in \bullet^t$ and produces $Post(t, p)$ tokens in the output places $p \in t^{\bullet}$. The next marking is $M' = M + Post(t, \cdot) - Pre(\cdot, t)$. A trace τ is defined as $\tau = M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots \xrightarrow{t_k} M_k$, where for $i = 1 \dots k$, $M_{i-1} \geq Pre(t_i)$. $M_0 \xrightarrow{\tau} M_k$ denotes that the sequence τ may fire at M_0 yielding M_k .

A PN $\langle \mathcal{N}, M_0 \rangle$ is *1-safe* if for every place $p \in \mathcal{P}$ we have that $M(p) \leq 1$ for any marking M that is reachable from M_0 .

2.2 Occurrence nets

Definition 1 Given a PN $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F)$ the immediate dependence relation $\preceq_1 \subset (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})$ is defined as:

$$\forall (a, b) \in (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P}) : a \preceq_1 b \text{ if } F(a, b) \neq 0$$

Define \preceq as the transitive closure of \preceq_1 ($\preceq = \preceq_1^*$).

The immediate conflict relation $\#_1 \subset \mathcal{T} \times \mathcal{T}$ is defined as:

$$\forall (t_1, t_2) \in \mathcal{T} \times \mathcal{T} : t_1 \#_1 t_2 \text{ if } \bullet^{t_1} \cap \bullet^{t_2} \neq \emptyset$$

Define $\# \subset (\mathcal{P} \cup \mathcal{T}) \times (\mathcal{P} \cup \mathcal{T})$ as $\forall (a, b) \in (\mathcal{P} \cup \mathcal{T}) \times (\mathcal{P} \cup \mathcal{T}) : a \# b$ if $\exists t_1, t_2$ s.t. $t_1 \#_1 t_2$ and $t_1 \preceq a$ and $t_2 \preceq b$.

The independence relation $\parallel \subset (\mathcal{P} \cup \mathcal{T}) \times (\mathcal{P} \cup \mathcal{T})$ is defined as $\forall (a, b) \in (\mathcal{P} \cup \mathcal{T}) \times (\mathcal{P} \cup \mathcal{T})$:

$$a \parallel b \Rightarrow \neg(a \# b) \wedge (a \not\preceq b) \wedge (b \not\preceq a)$$

Definition 2 Given two PNs $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F)$ and $\mathcal{N}' = (\mathcal{P}', \mathcal{T}', F')$, ϕ is a homomorphism from \mathcal{N} to \mathcal{N}' , denoted $\phi : \mathcal{N} \rightarrow \mathcal{N}'$ where i) $\phi(\mathcal{P}) \subseteq \mathcal{P}'$ and $\phi(\mathcal{T}) \subseteq \mathcal{T}'$ and ii) $\forall t \in \mathcal{T}$, the restriction of ϕ to \bullet^t respectively t^{\bullet} is a bijection between \bullet^t and $\bullet^{\phi(t)}$ respectively between t^{\bullet} and $\phi(t)^{\bullet}$.

Definition 3 An occurrence net is a net $O = (B, E, \preceq_1)$ such that: i) $\forall a \in B \cup E : \neg(a \preceq a)$ (acyclic); ii) $\forall a \in B \cup E : |\{b : a \preceq b\}| < \infty$ (well-formed); iii) $\forall b \in B : |\bullet^b| \leq 1$ (no backward conflict).

In the following B is referred as the set of conditions while E is the set of events.

Definition 4 A configuration $C = (B_C, E_C, \preceq)$ in the occurrence net O is defined as follows:

i) C is a proper sub-net of O ($C \subseteq O$)

ii) C is conflict free, i.e.

$$\forall a, b \in (B_C \cup E_C) \times (B_C \cup E_C) \Rightarrow \neg(a \# b)$$

iii) C is causally upward-closed, i.e.

$$\forall b \in B_C \cup E_C : a \in B \cup E \text{ and } a \preceq_1 b \Rightarrow a \in B_C \cup E_C$$

iv) $\min_{\preceq}(C) = \min_{\preceq}(O)$

Definition 5 Consider a PN $\langle \mathcal{N}, M_0 \rangle$ s.t. $\forall p \in \mathcal{P} : M_0(p) \in \{0, 1\}$. A branching process \mathcal{B} of a PN $\langle \mathcal{N}, M_0 \rangle$ is a pair $\mathcal{B} = (O, \phi)$ where O is an occurrence net and ϕ is a homomorphism $\phi : O \rightarrow \mathcal{N}$ s.t.:

1. the restriction of ϕ to $\min_{\preceq}(O)$ is a bijection between $\min_{\preceq}(O)$ and M_0 (the set of initially marked places)

2. $\phi(B) \subseteq \mathcal{P}$ and $\phi(E) \subseteq \mathcal{T}$
3. $\forall a, b \in E : (\bullet a = \bullet b) \wedge (\phi(a) = \phi(b)) \Rightarrow a = b$

For a configuration C in O denote by $CUT(C)$ the maximal (w.r.t. set inclusion) set of conditions in C that have no successors in C :

$$CUT(C) = ((\bigcup_{e \in E_C} e^\bullet) \cup (\min_{\preceq}(O)) \setminus (\bigcup_{e \in E_C} \bullet e))$$

Definition 6 Given a PN $\langle \mathcal{N}, M_0 \rangle$ and two branching processes $\mathcal{B}, \mathcal{B}'$ of PN $\langle \mathcal{N}, M_0 \rangle$ then $\mathcal{B}' \subseteq \mathcal{B}$ if there exists an injective homomorphism $\varphi : \mathcal{B}' \rightarrow \mathcal{B}$ s.t. $\varphi(\min(\mathcal{B}')) = \min(\mathcal{B})$ and $\phi \circ \varphi = \phi'$.

There exists a unique (up to an isomorphism) maximum branching process (w.r.t. \subseteq) that is the unfolding of $\langle \mathcal{N}, M_0 \rangle$ and is denoted $\mathcal{U}_{\mathcal{N}}(M_0)$ [8].

Denote by \mathcal{C} the set of all the configurations C of the occurrence net $\mathcal{U}_{\mathcal{N}}(M_0)$. For a configuration $C \in \mathcal{C}$ denote by $\langle E_C \rangle_{\preceq}$ the set of strings that are linearizations of (E_C, \preceq) where a string $\sigma = e_1 e_2 \dots e_v$ is a linearization of (E_C, \preceq) if $v = |E_C|$ and $\forall e_\iota, e_\lambda \in E_C$ we have that: *i*) $e_\iota = e_\lambda \Rightarrow \iota = \lambda$ and *ii*) for $\iota \neq \lambda$, if $e_\iota \preceq e_\lambda$ then $\iota < \lambda$.

2.3 Time Petri nets

A Time Petri Net (TPN) $\mathcal{N}^\theta = (\mathcal{P}, \mathcal{T}, F, I^s)$, consists of an (untimed) Petri Net $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F)$ (called the untimed support of \mathcal{N}^θ) and the static time interval function $I^s : \mathcal{T} \rightarrow \mathcal{I}(\mathbb{Q}^+)$, $I_t^s = [L_t^s, U_t^s]$, $L_t^s, U_t^s \in \mathbb{Q}^+$, representing the set of all possible time delays associated to transition $t \in \mathcal{T}$.

In a TPN $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ we say that a transition t becomes enabled at the time θ_t^{en} then the clock attached to t is started and the transition t can and must fire at some time $\theta_t \in [\theta_t^{en} + L_t^s, \theta_t^{en} + U_t^s]$, provided t did not become disabled because of the firing of another transition. Notice that t is forced to fire if it is still enabled at the time $\theta_t^{en} + U_t^s$.

Definition 7 A state at the time θ (according to a global clock) of a TPN $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ is a pair $S_\theta = (M, FI)$ where M is a marking and FI is a firing interval function associated with each enabled transition in M ($FI : \mathcal{T} \rightarrow \mathcal{I}(\mathbb{Q}^+)$).

If t is executed at the time $\theta_t \in \mathbb{Q}^+$ we write $(M, FI) \xrightarrow{\langle t, \theta_t \rangle} (M', FI')$ or simply $S \xrightarrow{\langle t, \theta_t \rangle} S'$ where:

1. $(M \geq Pre(\cdot, t) \wedge \theta_t \geq \theta_t^{en} + L_t^s) \wedge (\forall t' \in \mathcal{T} \text{ s.t. } M \geq Pre(\cdot, t') \Rightarrow \theta_t \leq \theta_{t'}^{en} + U_{t'}^s)$
2. $M' = M - Pre(\cdot, t) + Post(t, \cdot)$
3. $\forall t'' \in \mathcal{T} \text{ s.t. } M' \geq Pre(\cdot, t'')$ we have:
 - (a) if $t'' \neq t$ and $M \geq Pre(\cdot, t'')$ then $FI(t'') = [\max(\theta_{t''}^{en} + L_{t''}^s, \theta_t), \theta_{t''}^{en} + U_{t''}^s]$
 - (b) else $\theta_{t''}^{en} = \theta_t$ and $FI(t'') = [\theta_{t''}^{en} + L_{t''}^s, \theta_{t''}^{en} + U_{t''}^s]$

A legal time trace τ^θ in a TPN \mathcal{N}^θ satisfies: $\tau^\theta = S_0 \xrightarrow{\langle t_1, \theta_{t_1} \rangle} S_1 \xrightarrow{\langle t_2, \theta_{t_2} \rangle} \dots \xrightarrow{\langle t_v, \theta_{t_v} \rangle} S_v$ where $S_\nu \xrightarrow{\langle t_{\nu+1}, \theta_{t_{\nu+1}} \rangle} S_{\nu+1}$ for $\nu = 0, \dots, v-1$.

In the following for a time trace τ^θ we use the notation τ to denote its untimed support. For the initial state S_0 we use also the notation M_0^θ . Denote $\mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ the set of all legal time

traces that can be executed in $\langle \mathcal{N}^\theta, M_0^\theta \rangle$. We call $\mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ the time language of the TPN $\langle \mathcal{N}^\theta, M_0^\theta \rangle$.

$\mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ is the untimed support language of the time language $\mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ i.e. $\mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta) = \{\tau \mid \exists \tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)\}$.

3 Diagnosis of TPNs

We consider the following plant description:

1. the TPN model $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ is untimed 1-safe
2. $\mathcal{T} = \mathcal{T}_o \cup \mathcal{T}_{uo}$ where \mathcal{T}_o is the set of observable transitions and \mathcal{T}_{uo} is the set of unobservable transitions
3. l_o is the observation labeling function $l_o : \mathcal{T} \rightarrow \Omega_o \cup \{\epsilon\}$ where Ω_o is a set of labels and ϵ is the empty label. $l_o(t) = \epsilon$ if $t \in \mathcal{T}_{uo}$ and $l_o(t) \in \Omega_o$ if $t \in \mathcal{T}_o$
4. when an observable transition $t^o \in \mathcal{T}_o$ is executed in the plant the label $l_o(t^o)$ is emitted together with the global time $\theta_{l_o(t^o)}$ when this execution of t^o took place
5. the execution of an unobservable transition does not emit anything (is silent)
6. the faults are modeled by a subset of unobservable events, $\mathcal{T}_f \subseteq \mathcal{T}_{uo}$; $l_f : \mathcal{T}_{uo} \rightarrow \Omega_f \cup \{\epsilon\}$ is the fault labeling function (Ω_f is a set of labels and ϵ is the empty label); $l_f(t) = \epsilon$ if $t \in \mathcal{T}_{uo} \setminus \mathcal{T}_f$ and $l_f(t) \in \Omega_f$ if $t \in \mathcal{T}_f$
7. the faults are unpredictable, i.e. $\forall t \in \mathcal{T}_f, \exists t' \in \mathcal{T} \setminus \mathcal{T}_f$ s.t. *i*) $\bullet t' \subseteq \bullet t$ and *ii*) $L_{t'}^s \leq U_t^s$.

The plant observation at the time the n^{th} observed event is executed in the plant is denoted as $\mathcal{O}_n^\theta = \langle obs_1, \theta_{obs_1}, \dots, \langle obs_n, \theta_{obs_n} \rangle \rangle$, where $obs_1, \dots, obs_n \in \Omega_o$ are the labels that are received and $\theta_{obs_1} \leq \theta_{obs_2} \leq \dots \leq \theta_{obs_n}$ are the occurrence times of the corresponding events.

Denote by $\mathcal{O}_{n,\xi}^\theta$ the plant observation at the time $\xi > \theta_{obs_n}$, i.e. $\mathcal{O}_{n,\xi}^\theta$ includes also the information that no observable event occurred in the interval $[\theta_{obs_n}, \xi]$.

$\mathcal{L}_{\mathcal{N}^\theta}^\theta(\mathcal{O}_n^\theta)$ is the set of all time traces that are feasible in $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ up to the time of the last observation θ_{obs_n} and that obey the received observation \mathcal{O}_n^θ where $\tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(\mathcal{O}_n^\theta)$ if: *i*) $\tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta, \theta_{obs_n})$ (τ^θ is legal); *ii*) $l_o(\tau) = obs_1, \dots, obs_n$ (τ^θ obeys the "untimed" observation), and *iii*) for each observable transition $t_k^o \in \mathcal{T}_o$, $k = 1, \dots, n$ we have that $l_o(t_k^o) = obs_k \Rightarrow \theta_{t_k} = \theta_{obs_k}$ (τ^θ obeys the execution times of the observed transitions).

Similarly $\mathcal{L}_{\mathcal{N}^\theta}^\theta(\mathcal{O}_{n,\xi}^\theta)$ is the set of all time traces that are feasible in $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ up to the time ξ and that obey the received observation $\mathcal{O}_{n,\xi}^\theta$.

The plant diagnosis $\mathcal{D}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta)$ based on the received observation $\mathcal{O}_{n,\xi}^\theta$ comprises the untimed strings obtained by projecting the untimed support traces contained in $\mathcal{L}_{\mathcal{N}^\theta}^\theta(\mathcal{O}_{n,\xi}^\theta)$ onto the set of fault transitions \mathcal{T}_f :

$$\mathcal{D}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta) = \left\{ \tau_f \mid \tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(\mathcal{O}_{n,\xi}^\theta) \text{ and } \tau_f = l_f(\tau) \right\}$$

The diagnosis result $\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta)$ indicates that a fault for sure happened if all the traces contain fault events, i.e.

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta) = \{F\} \Leftrightarrow \epsilon \notin \mathcal{D}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta)$$

If $\mathcal{D}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta)$ contains only the empty string ϵ then the diagnosis result is normal, i.e.

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta) = \{\mathbb{N}\} \Leftrightarrow \mathcal{D}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta) = \{\epsilon\}.$$

Otherwise the diagnosis result is uncertain, i.e. a fault could have happened but did not necessarily happen [9].

4 The analysis based on partial orders

The partial order reduction techniques developed for untimed PN [8] are shown in [1],[3] to be applicable for TPN. Consider a configuration C in the unfolding $\mathcal{U}_{\mathcal{N}}(M_0)$ of the untimed PN support of a TPN. Then consider a valuation Θ of the execution times at which the events $e \in E_C$ in the configuration C are executed, that is for each $e \in E_C$ consider a time value $\theta_e \in \mathbb{T}$ (\mathbb{T} the time axis) at which e occurs and Θ is an $|E_C|$ -tuple representing the execution times of all the events $e \in E_C$.

An untimed configuration C with a valuation $\Theta \in \mathbb{T}^{|E_C|}$ of the execution time for its events is called a time configuration of the TPN. A time configuration is legal if there is a legal trace $\tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ in the TPN $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ whose untimed support τ is a linearization of the partial order relation of the events in the configuration (i.e. $\tau = \phi(\sigma)$ and $\sigma \in \langle E_C \rangle_{\preceq}$) while the execution time θ_t of every transition t considered in the trace τ^θ is identical with the valuation of the event e for which t is its image via ϕ .

Consider an untimed configuration $C \in \mathcal{C}$. The TPN $C^\theta = (B_C, E_C, \preceq, \min_{\preceq}(\mathcal{U}_{\mathcal{N}}), I^s)$ is obtained by attaching to each event $e \in E_C$ the static interval I_e^s that corresponds in the original TPN to transition t s.t. $\phi(e) = t$.

Denote by \tilde{K}_{C^θ} the following system of inequalities:

$$\tilde{K}_{C^\theta} \left\{ \begin{array}{l} \max_{e' \in \bullet\bullet e} (\theta_{e'}) + L_e^s \leq \theta_e \leq \max_{e' \in \bullet\bullet e} (\theta_{e'}) + U_e^s \quad \forall e \in E_C \end{array} \right. \quad (1)$$

where in (1) $\bullet\bullet e = \emptyset$ implies $\max_{e' \in \bullet\bullet e} (\theta_{e'}) = 0$.

Proposition 1 $\forall \tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta)$ we have that if $\tau = \phi(\sigma)$ and $\sigma \in \langle E_C \rangle_{\preceq}$, then Θ is a solution of \tilde{K}_{C^θ} , where $\Theta = (\theta_{t_1}, \dots, \theta_{t_{|E_C|}}) = (\theta_{e_1}, \dots, \theta_{e_{|E_C|}})$ with $\phi(e_i) = t_i$, $i = 1, \dots, |E_C|$.

Proof: The proof is straightforward. \square

Denote by $Sol(\tilde{K}_{C^\theta})$ the set of all solutions of \tilde{K}_{C^θ} . The $|E_C|$ -hyperbox $\tilde{\mathbf{I}}$ that circumscribes $Sol(\tilde{K}_{C^\theta})$ is easily obtained as: $\forall e \in E_C$, $\tilde{\mathbf{I}}(e) = [\tilde{L}(e), \tilde{U}(e)]$ with $\tilde{L}(e) = \max_{e' \in \bullet\bullet e} (\tilde{L}(e')) + L_e^s$ and $\tilde{U}(e) = \max_{e' \in \bullet\bullet e} (\tilde{U}(e')) + U_e^s$ where $\forall e \in E_C$ s.t. $\bullet\bullet e = \emptyset$, $\tilde{L}(e) = L_e^s$ and $\tilde{U}(e) = U_e^s$.

Example 1 Consider the TPN displayed in Fig. 1. Static intervals are attached to each transition. The observable transitions are t_4 , t_7 and t_{10} and they emit the same label. t_3 and t_9 are faulty transitions.

In Fig. 2 a part of the unfolding $\mathcal{U}_{\mathcal{N}}(M_0)$ is displayed where attached to each event $e \in E$ is the interval $\tilde{\mathbf{I}}(e)$.

We cannot claim yet that for $C \in \mathcal{C}$ there exists at least a legal time configuration that corresponds with C because for a general TPN the enabling of a transition does not guarantee that it eventually fires because some conflicting transition may be forced to fire beforehand.




Figure 1:




Figure 2:

Denote by \check{E}_C the set of conflicting events of a configuration $C \in \mathcal{C}$ where \check{E}_C comprises the events that could have been executed but are not included in E_C :

$$\check{E}_C = \{\check{e} \in E \setminus E_C \mid \bullet\check{e} \subseteq B_C\}$$

The characteristic system K_{C^θ} of configuration $C \in \mathcal{C}$ is obtained by adding to \tilde{K}_{C^θ} inequalities regarding the conflicting events:

$$K_{C^\theta} = \left\{ \begin{array}{l} \max_{e' \in \bullet\bullet e} (\theta_{e'}) + L_e^s \leq \theta_e \leq \max_{e' \in \bullet\bullet e} (\theta_{e'}) + U_e^s \quad \forall e \in E_C \\ \min_{e' \#_1 \check{e}} (\theta_{e'}) \leq \max_{e' \in \bullet\bullet \check{e}} (\theta_{e'}) + U_{\check{e}}^s \quad \forall \check{e} \in \check{E}_C \end{array} \right.$$

Proposition 2 Given an arbitrary time ξ we have that $\tau^\theta \in \mathcal{L}_{\mathcal{N}^\theta}^\theta(M_0^\theta, \xi)$ iff: i) $\tau = \phi(\sigma)$, $\sigma \in \langle E_C \rangle_{\preceq}$ and $C \in \mathcal{C}$; ii) Θ is a solution of K_{C^θ} , iii) $\forall e \in E_C \Rightarrow \theta_e \leq \xi$, and iv) $\forall e \in ENABLED(C)$, $\theta_e \geq \xi$.

Proof: \Rightarrow Since the PN is 1-safe we have that for any legal untimed trace τ there exists a unique configuration C s.t. $\tau \in \langle E_C \rangle_{\preceq}$. Condition 1, 3 and 4 are trivial and the proof that $\Theta = (\theta_1, \dots, \theta_n)$ is a solution of K_{C^θ} is simply by induction. \Leftarrow The proof is trivial. \square

The problem the we should answer next is: "Up to what time ξ to make the calculations for the on-line monitoring?"

There are different solutions to answer this question, depending on the computational capability, the plant behavior, and the requirements for the diagnosis result.

Solution 1: Calculations in advance

The first solution is appropriate for a plant known to have a cyclic operation, where each operation cycle is initiated by the plant operator.

Having derived the plant behavior up to the time $\hat{\xi}$ that corresponds with the completion of an operation cycle, the plant is monitored on-line in the following way:

1. the received observation is taken in to account adding (in)equality constraints to the characteristic system of a configuration.
2. or configurations are discarded when the current time exceeds the latest execution time of an observable event in a configuration.

The main drawback of this method is that a large amount of calculations is performed in advance and then discarded because of the received observation.

Solution 2: Calculations after each observation

The second solution is to perform calculations each time an event is observed in the plant. E.g. when the first observable event is executed in the plant we derive the plant behavior up to the time θ_{obs_1} in the following way.

Let the first observation be $\mathcal{O}_1^\theta = \langle obs_1, \theta_{obs_1} \rangle$. Consider the set of configurations $\mathcal{C}(\mathcal{O}_1^\theta)$ s.t. $C \in \mathcal{C}(\mathcal{O}_1^\theta)$ if:

1. E_C contains only one event e^o s.t. $\phi(e^o) \in \mathcal{T}_o$ and $l_o(\phi(e^o)) = obs_1$, and $\theta_{obs_1} \in \tilde{I}(e^o)$
2. $\forall e \in \bullet CUT(C) : \tilde{L}(e) < \theta_{obs_1}$
3. $\forall e \in ENABLED(C) : \tilde{U}(e) > \theta_{obs_1}$

where $ENABLED(C)$ denotes the set of events that correspond to transitions that are enabled from $\phi(CUT(C))$.

The characteristic system $K_{C^\theta}(\mathcal{O}_1^\theta)$ of configuration $C^\theta \in \mathcal{C}(\mathcal{O}_1^\theta)$ is obtained by adding to \tilde{K}_{C^θ} inequalities regarding the conflicting events and the received observation.

This method requires less computation but the price to be paid is that a fault may be detected with a delay. This is because no calculations are performed until a new observation is received, thus the fact that the current time of the plant exceeds the latest execution time of an observable event is not taken in to account.

However this method is practically useful when the frequency of observations is high, i.e. the time interval in between two observations is short and control actions are inevitably taken with some latency. Moreover this method is also suitable when the plant observation is known to be uncertain, i.e. the observation of an event can be lost because of a sensor failure. This is because in between two observations the diagnosis result w.r.t. the detection of the faults that for sure happened does not change if the observation is uncertain.

Solution 3: Calculations up to a discarding time

A discarding time is the earliest time when in absence of any observation one can discard untimed support traces because it can be proved that they are not valid. E.g. the first discarding time is the smallest latest execution time of an observable transition in the plant.

Definition 8 A configuration $C_\nu \in \mathcal{C}$ is derived up to the time ξ if: i) $\max_{e \in \bullet CUT(C_\nu)}(\tilde{L}_\nu(e)) \leq \xi$ and ii) $\min_{e \in ENABLED(C_\nu)}(\tilde{U}_\nu(e)) > \xi$. Given a configuration $C_\nu \in \mathcal{C}$ that is derived up to a time ξ' , denote by $C_\nu(\xi)$ the set of extensions of C_ν up to the time $\xi > \xi'$ where $C_{\ell_\nu} \in C_\nu(\xi)$ if: i) $C_\nu \subseteq C_{\ell_\nu}$ (C_{ℓ_ν} is a continuation of C_ν) and ii) C_{ℓ_ν} is derived up to the time ξ .

The first discarding time $\hat{\theta}$ is calculated iteratively as follows. $\hat{\theta}$ is initiated with a big value (say $+\infty$ for simplicity) and then starting from the initial configuration $C^\perp = (B^\perp, E^\perp, \leq_1)$ we construct an initial part of the net unfolding by appending events as in the untimed case, the only difference being that among all the enabled events denoted by $ENABLED(C)$ only the events with the smallest upper limit $\tilde{U}(e)$ are appended, until the first observable event say e^o is encountered.

The discarding time is set equal to $\tilde{U}(e^o)$ and then the configurations that contain e^o are extended up to the time $\tilde{U}(e^o)$. Denote this set by \mathcal{C}_{obs}^{new} . Then for each configuration $C_\nu \in \mathcal{C}_{obs}^{new}$ we calculate $Sol(K_{C_\nu^\theta})$ and for those configurations that have a non-empty solution set we calculate $U_\nu(e^o)$, i.e. the smallest latest time when an observable event e^o can be executed. Obviously $U_\nu(e^o) \leq \tilde{U}_\nu(e^o)$.

The discarding time $\hat{\theta}$ is set as the smallest latest time when an observable event can be forced to execute considering all $C_\nu \in \mathcal{C}_{obs}^{new}$. Notice that a configuration C_ν may contain some other observable events and after calculating $Sol(K_{C_\nu^\theta})$ some other observable event may have the smallest latest time for its execution. Then recursively all the configurations that contain only unobservable events are extended up to the new discarding time $\hat{\theta}$ by appending event(s) selected among all the enabled with the smallest upper limit $\tilde{U}(e)$. Continue this operation until either a new observable event is encountered or no more events can be appended.

Notice that because $\hat{\theta}$ is calculated recursively some configurations (that contain at least one observable events) are derived up to times bigger than $\hat{\theta}$. However this does not affect the diagnosis result since the events that can be executed after the time $\hat{\theta}$ are seen as a prognosis.

The on-line diagnosis algorithm works as follows. When the process starts we derive the set of configurations up to the first discarding time and then we have two cases:

Case 1 If no observation is received before the time $\hat{\theta}$ then:

1. the configurations that contain observable events having the upper limit equal to $\hat{\theta}$ are discarded
2. for all the other configurations that contain observable events inequalities of the form:

$$\mathcal{K}_{obs} = \left\{ \theta_{e^o} > \hat{\theta} \mid e^o \in E_{C_\nu} \text{ and } \phi(e^o) \in \mathcal{T}_o \right\}$$

are added to the characteristic systems $K_{C_\nu^\theta}$ and we derive the entire solution set

3. for all the configurations $C_\nu \in \mathcal{C}_{uno}$ that contain only unobservable events we check only if $Sol(K_{C_\nu^\theta})$ has a non-empty set of solutions.
4. denote by $\mathcal{E}(\mathcal{O}_{0,\hat{\theta}}^\theta)$ the set of traces that are obtained as linearizations of the set of events of the configurations that are not discarded.
5. the diagnosis $\mathcal{D}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_{0,\hat{\theta}}^\theta)$ is obtained projecting $\mathcal{E}(\mathcal{O}_{0,\hat{\theta}}^\theta)$ onto the set of fault transitions \mathcal{T}_f

Case 2 If the first observation $\langle obs_1, \theta_{obs_1} \rangle$ is received before the time of the process becomes $\hat{\theta}$ then:

1. the set of configurations \mathcal{C}_{unu} that contain only unobservable events is discarded
2. for each configuration $C_\nu \in \mathcal{C}_{obs}$ that contains observable events an equality relation:

$$\mathcal{K}'_{obs_1} = \{\theta_{e^o} = \theta_{obs_1} \mid l_o(e^o) = obs_1 \wedge e^o \in C_\nu\}$$

and for observable events other than e^o inequalities of form:

$$\mathcal{K}''_{obs_1} = \{\theta_{e'^o} > \hat{\theta} \mid e'^o \in E_C \text{ and } \phi(e'^o) \in \mathcal{T}_f\}$$

are added to the characteristic system K_{C_ν} and then we derive the entire solution set

3. denote by $\mathcal{E}(\mathcal{O}_1^\theta)$ the set of traces that are obtained as linearizations of the set of events of the configurations that are not discarded.
4. $\mathcal{D}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_1^\theta)$ is obtained projecting $\mathcal{E}(\mathcal{O}_1^\theta)$ onto \mathcal{T}_f

Notice that the plant diagnosis is derived either at the time of the first observed event $\mathcal{D}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_1^\theta)$ or in absence of any observation at the first discarding time $\hat{\theta}$, $\mathcal{D}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_{0,\hat{\theta}}^\theta)$.

Theorem 1 *Given a TPN model $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ we have that:*

1. when the first observable event is executed:

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_1^\theta) = \{F\} \Leftrightarrow \mathcal{DR}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_1^\theta) = \{F\}$$

2. if no observation is received until the first discarding $\hat{\theta}$:

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{0,\hat{\theta}}^\theta) = \{F\} \Leftrightarrow \mathcal{DR}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_{0,\hat{\theta}}^\theta) = \{F\}$$

3. and for any time $\xi \leq \hat{\theta}$, in absence of any observation, the diagnosis result is different from F :

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{0,\xi}^\theta) \neq \{F\}$$

Proof: (1) and (2) have a similar proof. Based on Proposition 2 we calculate the set of legal traces up to given time ξ . However some configurations include events that are executed after the time θ_{obs_1} or $\hat{\theta}$. Since the faults are unpredictable the consideration of some events that can be executed after the time θ_{obs_1} or $\hat{\theta}$ does not change the diagnosis result w.r.t. the detection of faults that for sure happened. (3) is proved straightforwardly by the assumption that the faults are unpredictable. \square

Remark 1 *Obviously by imposing the inequality that all the events in a configuration have execution times smaller than θ_{obs_1} or $\hat{\theta}$ allows one to derive exactly the diagnosis result by removing the events that can be executed after the time θ_{obs_1} respectively $\hat{\theta}$. However this is not efficient for practical calculations especially when the frequency of observations is high. Notice also that calculations in advance are not fully developed, thus it may be that an event that is considered executed after θ_{obs_1} might not be executed since an event that is successor of the observed event can pre-empt its execution.*

In what follows we present two methods to derive the solution set of the characteristic system of a configuration. The first method is based on the ELCP and derives the entire solution set as a union of faces of a polyhedron that satisfy the cross-complementarity condition [10].

The second method is based on constraint propagation and derives for a configuration C a set of $|E_C|$ -hyperboxes s.t. the union of the subsets of solutions that are circumscribed by the $|E_C|$ -hyperboxes is a cover of the entire solution set.

5 The method based on ELCP

The ELCP is defined as follows (see [10]). Given $A \in \mathbb{R}^{w \times z}$, $G \in \mathbb{R}^{q \times z}$, $c \in \mathbb{R}^w$, $d \in \mathbb{R}^q$, and m index sets $\psi_1, \dots, \psi_m \subseteq \{1, \dots, w\}$, find $x \in \mathbb{R}^z$ such that

$$Ax \geq c, \quad Gx = d \quad (2)$$

$$\sum_{j=1}^m \prod_{i \in \psi_j} (Ax - c)_i = 0. \quad (3)$$

Condition (3) can be interpreted as follows. Since $Ax \geq c$, all the terms in (3) are nonnegative. Hence, (3) is equivalent to $\prod_{i \in \psi_j} (Ax - c)_i = 0$ for $j = 1, \dots, m$. So we could say that each set ψ_j corresponds to a group of inequalities in $Ax \geq c$, and that in each group at least one inequality should hold with equality. In [10] we have developed an algorithm to find *all* solutions of an ELCP. This algorithm yields a description of the complete solution set of an ELCP by finite points, generators for extreme rays, and a basis for the linear subspace associated with the maximal affine subspace of the solution set of the ELCP.

Let us now explain how $(max, +)$ equations of the form

$$\max_{i \in \mathcal{J}} (\theta_i) + L \leq \theta \leq \max_{i \in \mathcal{J}} (\theta_i) + U \quad (4)$$

can be recast as an ELCP. First we introduce a dummy variable $\gamma = \max_{i \in \mathcal{J}} \theta_i$. Then (4) reduces to

$$\gamma + L \leq \theta \leq \gamma + U, \quad (5)$$

which already fits the ELCP format. Let us now look at the equation $\gamma = \max_{i \in \mathcal{J}} \theta_i$. This can be recast as

$$\gamma \geq \theta_i \quad \text{for all } i \in \mathcal{J}, \quad (6)$$

where for at least one index $i \in \mathcal{J}$ equality should hold, i.e.,

$$\prod_{i \in \mathcal{J}} (\gamma - \theta_i) = 0. \quad (7)$$

Clearly, equations (5)–(7) constitute an ELCP.

Thus K_{C^θ} can be treated as an ELCP. First we derive the polyhedron that provides the set of solution for the system of linear (in)equalities given by 2. The solution set of the ELCP is obtained as a union of faces of a polyhedron that satisfy the cross-complementarity condition [10].

6 The method based on constraint propagation

Before formally presenting the second algorithm we introduce first the definition of a time interval configuration.

A time interval configuration $C(\mathbf{I})$ is an untimed configuration $C \in \mathcal{C}$ endowed with time intervals for the execution of the events within the configuration. \mathbf{I} is a vector of dimension $|E_C|$ that comprises for each event $e \in E_C$ the time interval $I(e)$ in which the event e is assumed executed.

Definition 9 *Given the observation \mathcal{O}_1^θ and a configuration $C \in \mathcal{C}(\mathcal{O}_1^\theta)$ we have that the time interval configuration $C(\mathbf{I})$ is legal if for any event e_i ($\forall e_i \in E_C$) and for any execution time θ_{e_i} of the event e_i ($\forall \theta_{e_i} \in I(e_i)$) there exist execution times for all the other events within the configuration ($\exists \theta_{e_j} \in I(e_j)$ for all $e_j \in E_C \setminus \{e_i\}$) s.t. $\Theta = (\theta_{e_1}, \dots, \theta_{e_i}, \dots, \theta_{e_{|E_C|}})$ is a solution of the characteristic system K_{C^θ} ($\Theta \in \text{Sol}(K_{C^\theta})$).*

Given a hyperbox $\mathbf{I}_\nu \subseteq \mathbf{I}$ denote by $[L_\nu(e), U_\nu(e)]$ the execution time interval for the event e . Then for a conflicting event \check{e} denote by $L_\nu(\check{e}) = \max_{e' \in \bullet\bullet_{\check{e}}}(L_\nu(e')) + U_\check{e}^s$ and $U_\nu(\check{e}) = \max_{e' \in \bullet\bullet_{\check{e}}}(U_\nu(e')) + U_\check{e}^s$ the earliest respectively the latest time when \check{e} is forced to fire. We have that.

Proposition 3 $C(\mathbf{I}_\nu)$ is a legal time interval configuration if the following conditions hold true:

1. $\mathbf{I}_\nu \subseteq \tilde{\mathbf{I}}$ such that $L_\nu(e) \leq \max_{e' \in \bullet\bullet_e}(L_\nu(e')) + U_e^s$ and $U_\nu(e) \geq \max_{e' \in \bullet\bullet_e}(U_\nu(e')) + L_e^s$
2. $\forall \check{e} \in \check{E}_C, \exists e \in E_C$ s.t. $e \#_1 \check{e}$ and $L_\nu(e) \leq \check{L}_\nu(\check{e})$ and $U_\nu(e) \leq \check{U}_\nu(\check{e})$.
3. $\theta_{obs_1} = \theta_{e^o}$ for $e^o \in E_C$, $\phi(e^o) = l(obs_1)$
4. $\forall e \in \bullet CUT(C) \Rightarrow U_\nu(e) \leq \theta_{obs_1}$
5. $\forall e \in ENABLED(C) \Rightarrow \max_{e' \in \bullet\bullet_e}(L_\nu(e')) + U_e^s \geq \theta_{obs_1}$.

Proof: The proof is lengthy and is omitted. \square

In the following we present an algorithm that derives a set of $|E_C|$ -hyperboxes, $\{\mathbf{I}_\nu \mid \nu \in \mathcal{V}\}$ (\mathcal{V} the set of indexes) s.t. for each $|E_C|$ -hyperbox \mathbf{I}_ν , $C(\mathbf{I}_\nu)$ is a legal time interval configuration and the union of the subsets $\{Sol_\nu(K_{C^\theta}) \mid \nu \in \mathcal{V}\}$ that are circumscribed by \mathbf{I}_ν is a cover of the entire solution set $Sol(K_{C^\theta})$, i.e. $\bigcup_{\nu \in \mathcal{V}} Sol_\nu(K_{C^\theta}) = Sol(K_{C^\theta})$, where $Sol_\nu(K_{C^\theta}) = Sol(K_{C^\theta}) \cap \mathbf{I}_\nu$.

The idea behind developing the algorithm that we propose is as follows. First we calculate the hyperbox $\tilde{\mathbf{I}}$ that circumscribes $Sol(K_{C^\theta})$. Then we should impose the timing constraints imposed by the conditions 2 – 5 in Proposition 3. We have three kinds of constraints. Denote by \mathcal{K}_{conf} , \mathcal{K}'_{obs} , and \mathcal{K}''_{obs} the set of constraints imposed by the set of conflicting events (condition (2)), the equality constraint required by the observation of the label l_{obs_1} (condition (3)), and respectively the set of constraints that require that the time configuration is complete w.r.t. the time θ_{obs_1} (none of the concurrent parts of the process are left behind in time).

Consider a constraint κ_e on the time interval $\tilde{I}(e) = [\tilde{L}(e), \tilde{U}(e)]$ of an event $e \in E_C$ where:

$$\kappa_e := \left\{ I'(e) = [L'(e), U'(e)] \mid L'(e) > \tilde{L}(e) \text{ or } U'(e) < \tilde{U}(e) \right\}$$

The set of solutions of \tilde{K}_{C^θ} that satisfy κ_e , denoted $Sol(\tilde{K}_{C^\theta} \wedge \kappa_e)$, is obtained propagating the constraint κ_e forward to its successors and backwards to its predecessors:

- *forward propagation:* for all $e_v \in e^{\bullet\bullet}$:

$$L'(e_v) = \max(\tilde{L}(e) + L_{e_v}^s, \tilde{L}(e_v)) \text{ and}$$

$$U'(e_v) = \min(\tilde{U}(e) + U_{e_v}^s, \tilde{U}(e_v))$$
- *backward propagation:*
 - i) for all $e_v \in \bullet\bullet_e$:

$$U'(e_v) = \min(\tilde{U}(e) - L_e^s, \tilde{U}(e_v))$$
 - ii) for each $e_v \in \bullet\bullet_e$ s.t. $\tilde{L}(e) - U_e^s > \tilde{U}(e_v)$ consider a different case $\nu \in \mathcal{V}$:
 - ii.1) $L'_\nu(e_v) = \tilde{L}(e) - U_e^s$
 - ii.2) for all $e_i \in \bullet\bullet_{e_v}$, $e_i \neq e_v$: $L'_\nu(e_i) = \tilde{L}(e_i)$.




Figure 3:

The backward propagation of a constraint κ_e may require to split an $|E_C|$ -hyperbox considering different cases. Notice that the number of cases is not bigger than the number of concurrent predecessor events of the event e to whom the constraint κ_e is applied. For each hyperbox $\mathbf{I}_{\nu'}$, $\nu' \in \mathcal{V}'$ the set of constraints is updated since in general it may be that new constraints appear while some of the previous constraints are satisfied. If a constraint cannot be imposed the case is aborted while if the set of constraints is empty the algorithm returns an hyperbox that circumscribes a subset of solutions of K_{C^θ} .

The constraint propagation algorithm works as follows:

1. first step is to impose the constraints of kind \mathcal{K}'_{obs} and \mathcal{K}''_{obs} (required by the received observation)
2. the second step is to impose for each $|E_C|$ -hyperbox that results after step 1, the set of constraints \mathcal{K}_{conf} . E.g. for \mathbf{I}_ν consider that $\exists \check{e} \in E_C$ s.t. condition 2 in Proposition 3 is not satisfied. Then for each $e \in E_C$ s.t. $e \#_1 \check{e}$ consider a different case and impose a constraint $\kappa_e := \{L'_{\nu'}(e) = L_{\nu'}(\check{e})\}$ if $L_{\nu'}(\check{e}) \leq L_{\nu'}(e)$ or $\kappa_{\check{e}} = \{U'_{\nu'}(\check{e}) = U_{\nu'}(e)\}$ if $U_{\nu'}(\check{e}) \leq U_{\nu'}(e)$.
3. an arbitrary constraint κ_e or $\kappa_{\check{e}}$ is selected and then it is imposed backwards. If new constraints appear on the time intervals of the predecessor events of e or \check{e} then one of these constraints is selected and it is imposed further backwards until a decision is achieved. Then constraints are propagated forward for the $|E_C|$ -hyperboxes that are not aborted. The maximum number of different cases that result propagating recursively a constraint backwards is smaller than the size of maximum set of concurrent events in the configuration
4. a decision is achieved for each case in finite time since the corner points of each $|E_C|$ -hyperbox are rational numbers and each constraint that is applied either reduces one edge of the $|E_C|$ -hyperbox or returns success/abort.

Example 2 Consider for the configuration C displayed in Fig. 3 that the first observation is received at the time 13 and consider the case when e_4 is the event that was observed. Let $\kappa'_{e_4} = \{\theta_{e_4} = 13\}$. κ'_{e_4} is propagated backwards and a new constraint κ'_{e_3} appears where $\kappa'_{e_3} = \{I_{e_4} = [5, 9]\}$. κ'_{e_3} is propagated backwards but no new constraints appears. Then e_{10} is required to be executed after $\theta_{e_4} = 13$, i.e. $\kappa''_{e_{10}} = \{\theta_{10} \in [13, 17]\}$. $\kappa''_{e_{10}}$ is propagated backwards and

a constraint κ_{e_9} appears where $\kappa_{e_9}'' = \{I_{e_9} = [4, 8]\}$. κ_{e_9}'' is propagated backwards and no new constraint appears.

Then the timing constraints required by the conflicting events \check{e}_1 and \check{e}_{12} are satisfied. What is left is the conflicting event \check{e}_6 . We have that $e_3 \# \check{e}_6$ and $e_9 \# \check{e}_6$ and $I(e_3) = [5, 9]$, $I(e_9) = [4, 8]$, and $I(\check{e}_6) = [3, 7]$.

We have two cases. First consider $e_3 \# \check{e}_6$. We have $\kappa_{\check{e}_6} = \{L'_{e_6} = 5\}$ and $\kappa_{e_3} = \{U'_{e_3} = 7\}$. $\kappa_{\check{e}_6}$ is propagated backwards and we have two cases: either $I_1(e_5) = [2, 5]$ and $I_1(e_8) = [1, 4]$ or $I_2(e_5) = [1, 5]$ and $I_2(e_8) = [2, 4]$. $\kappa_{e_3} = \{U'_{e_3} = 7\}$ does not produce new constraints. We obtain two hyperboxes and if we consider the case when $e_9 \# \check{e}_6$ we obtain in a similar way another two hyperboxes.

7 The on-line diagnosis

In the previous sections we have presented the plant diagnosis up to the first observation or in absence of any observation up to the first discarding time. Then the on-line diagnosis is performed calculating the plant behavior up to a new discarding time.

Theorem 2 Given a TPN model $\langle \mathcal{N}^\theta, M_0^\theta \rangle$ we have that:

1. when an observable event is executed:

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_n^\theta) = \{F\} \Leftrightarrow \mathcal{DR}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_n^\theta) = \{F\}$$
2. for $\hat{\theta}$ the first discarding time after the time when the n^{th} observed event is reported:

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\hat{\theta}}^\theta) = \{F\} \Leftrightarrow \mathcal{DR}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_{n,\hat{\theta}}^\theta) = \{F\}$$
3. and in absence of any observation, the diagnosis result w.r.t. the detection of the faults that for sure happened calculated any time in between the last observed event and the discarding time is constant, i.e. $\forall \xi \in [\theta_{obs_n}, \hat{\theta})$:

$$\mathcal{DR}_{\mathcal{N}^\theta}(\mathcal{O}_{n,\xi}^\theta) = \{F\} \Leftrightarrow \mathcal{DR}_{\mathcal{N}^\theta}^{po}(\mathcal{O}_n^\theta) = \{F\}.$$

Proof: The proof is similar to the proof of Theorem 1. \square

8 Final remarks and future work

We have derived in this paper on-line algorithms for the diagnosis of TPN models. The plant behavior is derived up to a discarding time, i.e. up to a time when in absence of any observation one can discard untimed support traces because they are not consistent with the plant behavior. The analysis is based on partial orders and it requires to derive the solution set of systems of $(\max, +)$ -linear inequalities.

We have presented two algorithms to derive the entire solution set, one based on the ELCP and the second one based on constraint propagation. Both algorithms are NP-hard problems. Beside the number of events, the number of conflicting events, and the maximum number of predecessors respectively successors of a node in a configuration, the computational complexity of both methods strongly depends on the structure of the system.

However there are a few reasons that allow us to claim that the two methods are computationally more efficient than the ones ([1], [5]) presented in the literature. Comparing with the method based on the state class graph computation [5] our methods have the advantage that not all the interleaving of the concurrent events are considered. Moreover the

computational complexity depends in our case on the size of the largest subnet that contains unobservable transitions whereas the computation complexity in [5] depends on the size of the entire net. The algorithm in [1] solves a system of $(\max, +)$ -inequalities enumerating all the cases for each max-term. This combinatorial approach is known in the literature to be computationally less efficient than the ELCP.

Finally notice that for the above example the ELCP provides 8 subsets while constraint satisfaction only finds 4 subsets. The reason is that each face of a polyhedron that satisfies a cross-complementarity condition provides a legal time interval configuration but the converse is not true. The subset of solutions that is circumscribed by the hyperbox of a time interval configuration may be obtained as a union of faces of a polyhedron that satisfy a cross-complementarity condition.

However the set of hyperboxes obtained running the algorithm based on constraint propagation does not allow one to calculate the minimum and maximum time separation between the execution of two events unless a further refinement of the calculations is performed.

We plan to extend the methodology for a distributed setting where the strong assumptions considered in [6] to be relaxed.

References

- [1] T. Aura and J. Lilius. Time processes of Time Petri Nets. *ATPN'97 - LNCS*, 1248, 1997.
- [2] B. Berthomieu and M. Menasche. An enumerative approach for analyzing Time Petri Nets. *IFIP Congress, Paris*, 1983.
- [3] T. Chatain and C. Jard. Time supervision of concurrent systems using symbolic unfoldings of Time Petri Nets. *Int. Conf. on Formal Modeling and Analysis of Time Systems*, Uppsala, Sweden, 2005.
- [4] E. Fabre, A. Benvensite, S. Haar, and C. Jard. Distributed monitoring of concurrent and asynchronous systems. *Journal of Discrete Event Dynamic Systems*, 15(1):33–84, March 2005.
- [5] M. Ghazel, M. Bigand, and A. Toguyéni. A temporal-constraint based approach for monitoring of Discrete Event Systems under partial observation. In *IFAC Congress*, Prague, 2005.
- [6] G. Jiroveanu. *Fault diagnosis for large Petri Nets*. PhD thesis, Ghent University, Gent, Belgium, 2006.
- [7] G. Jiroveanu, B. De Schutter, and R.K. Boel. Fault Diagnosis for Time Petri Nets. In *Workshop on Discrete Event Systems (WODES'06)*, Ann Arbor, USA, 2006.
- [8] K. L. McMillan. Using unfoldings to avoid the state space explosion problem in verification of asynchronous circuits. In *4th Int. Workshop on CAV*, 1992.
- [9] M. Sampath, R. Sengupta, S. Lafortune, S. Sinnamo-hideen, and D. Teneketzis. Diagnosability of Discrete Event Systems. *IEEE-T on AC*, 40(9), 1995.
- [10] B. De Schutter and B. De Moor. The Extended Linear Complementarity Problem. *Mathematical Programming*, 71(3):289–325, Dec. 1995.