Effective Verification of Weak Diagnosability

Anoopam Agarwal, * **Agnes Madalinski,** ** **Stefan Haar** ***

* *Student, IIT Delhi, India (e-mail: anoopamagarwal@gmail.com)*
** *University Austral de Chile, Campus Miraflores, Valdivia, Chile (e-mail: amadalin@uach.cl)*
*** *INRIA and LSV (CNRS and ENS Cachan) 61, avenue du Président Wilson 94235 CACHAN Cedex, France (e-mail: haar@lsv.ens-cachan.fr,stefan.haar@inria.fr).*

Abstract: The *diagnosability problem* can be stated as follows: does a given labeled Discrete Event System allow for an outside observer to determine the occurrence of the "invisible" fault, no later than a bounded number of events after that unobservable occurrence, and based on the partial observation of the behaviour? When this problem is investigated in the context of concurrent systems, partial order semantics induces a separation between classical or strong diagnosability on the one hand, and *weak diagnosability* on the other hand. The present paper presents the first solution for checking weak diagnosability, via a *verifier* construction.

Keywords: Discrete event systems, diagnosis, Petri nets, events, observability, partial order semantics, Event structures.

1. INTRODUCTION

Diagnosis under partial observation is a classical problem in automatic control in general, and has received considerable attention in *discret event system (DES)* theory, among other fields. In the DES setting, the approach that we will call "classical" here supposes that the observed system is an automaton with transition set T, prefix-closed language $\mathcal{L} \subseteq T^*$, and a set of *observable transition labels* $\mathbb{O}$. The associated labeling map $\lambda : T \to \mathbb{O}$ may not be required injective, and leaves some transitions from T unobservable, in particular *fault* ϕ. The observations have the form of words $w \in \mathbb{O}^*$ obtained by applying λ to words in T^*. A classical definition of diagnosability is given in Sampath et al. (1995); we follow the equivalent presentaion of Cassandras and Lafortune (1999). Write $s \sim_\lambda s'$ iff $\lambda(s) = \lambda(s')$, and call any sequence s such that ϕ occurs in s a *faulty* sequence, and all other sequences *healthy*. Then :

Definition 1. (Sequential Diagnosability). A prefix-closed language $\mathcal{L} \subseteq T^*$ is *not (strongly) diagnosable* iff there exist sequences $s_N, s_Y \in \mathcal{L}$ such that:

(1) s_Y is faulty, s_N is healthy, and $s_N \sim_\lambda s_Y$;
(2) s_Y with the property 1 can be chosen arbitrarily long after the first fault, i. e. for every $k \in \mathbb{N}$ there exists a choice of $s_N, s_Y \in \mathcal{L}$ with the above properties and such that the suffix $s_{Y/\phi}$ of s_Y after the first occurrence of fault ϕ in s_Y satisfies $|s_Y| \geq k$.

Concurrent systems are difficult to supervise using the classical approach because of the state explosion problem. For intrinsically asynchronous distributed systems, such as encountered in telecommunications or more generally in networked systems, the use of models that reflect the local and distributed nature of the observed system, such as Petri nets or graph grammars, is helpful not only in terms of computational efficiency, but also *conceptually*. Putting these ideas together, Benveniste et al. (2003) extends diagnosis to asynchronous models *and their non-interleaved semantics*. This generalized methodology for fault diagnosis is based on the non-sequential executions of labeled Petri nets, that is, the partial order semantics in occurrence nets and event structures. Theoretical aspects of *partial order diagnosability* for Petri nets, in the spirit of the above definition, have been developped in Haar et al. (2003); Haar (2007, 2009, 2010a,b). While the sequential case is embedded and generalized in these results, new features emerge in partial ordered runs that have no counterpart in sequential behaviour; this led to the distinction between *strong* and *weak* observability and diagnosability properties in Haar et al. (2003); Haar (2010a).

2. PETRI NETS AND UNFOLDINGS

Definition 2. A **net** is a tuple $N = (P, T, F)$ where

- $P \neq \emptyset$ is a set of **places**,
- $T \neq \emptyset$ is a set of **transitions** such that $P \cap T = \emptyset$,
- $F \subseteq (P \times T) \cup (T \times P)$ is a set of **flow arcs**.

A **marking** is a multiset M of places, i.e. a map from P to N. A **Petri net** is a tuple $\mathcal{N} = (P, T, F, M)$, where *(i)* (P, T, F) is a finite net, and *(ii)* $M : P \to \mathbb{N}$ is an **initial marking**.

Elements of $P \cup T$ are called the *nodes* of $\mathcal{N}$. For a transition $t \in T$, we call ${}^\bullet t = \{p \mid (p, t) \in F\}$ the *preset*

of t, $t^\bullet = \{p \mid (t,p) \in F\}$ the *postset* of t. In Figure 1, we represent as usual places by empty circles, transitions by squares, F by arrows, and the marking of a place p by putting the corresponding number of *black tokens* into p. A transition t is *enabled* in marking M, written $M \xrightarrow{t}$ if $\forall p \in {}^\bullet t,\ M(p) > 0$; otherwise write $M \not\xrightarrow{t}$. An enabled transition t can *fire*, resulting in a new marking $m' = m - {}^\bullet t + t^\bullet$; this firing relation is denoted by $M \xrightarrow{t} M'$. A marking m is *reachable* if there exists a *firing sequence*, i.e. transitions $t_0 \dots t_n$ and markings $M_1, \dots, M_{n-1}$ such that $M_0 \xrightarrow{t_0} M_1 \xrightarrow{t_1} \dots \xrightarrow{t_n} M_n$. A Petri net is *safe* if for all reachable markings M, $M(p) \subseteq \{0,1\}$ for all $p \in P$; all Petri nets considered here are safe.

Occurrence nets and Unfoldings. In a net $N = (P, T, F)$, let $<_N$ the transitive closure of F, and $\leqslant_N$ the reflexive closure of $<_N$. For $t_1, t_2 \in T$, set $t_1 \#_{im} t_2$ and t_2 iff $t_1 \neq t_2$ and ${}^\bullet t_1 \cap {}^\bullet t_2 \neq \emptyset$, and define $\# = \#_N$ by

$$a \# b \Leftrightarrow \exists t_a, t_b \in T : \begin{cases} t_a \ \#_{im}\ t_b \\ \wedge\ t_a \ \leqslant_N\ a \\ \wedge\ t_b \ \leqslant_N\ b. \end{cases}$$

Define $\mathbf{co} = \mathbf{co}_{\mathcal{N}}$ by setting, for any $a, b \in P \cup T$,

$$a\ \mathbf{co}\ b \iff \neg(a \leqslant b)\ \wedge\ \neg(a \# b)\ \wedge\ \neg(b < a)$$

Definition 3. Let $ON = (B, E, G)$ be a net, and define the *closed* and *open* prime configurations for $e \in E$ by

$$[e] = \{y \in B \cup E \mid y \leqslant_{ON} e\}$$
$$\langle e \rangle \triangleq [e] \backslash \{e\}.$$

Then ON is an **occurrence net** if and only if it satisfies

(1) $\leqslant_{ON}$ is a partial order;
(2) for all $b \in B$, $|{}^\bullet b| \in \{0,1\}$;
(3) for all $e \in E$, the set $[e]$ is finite;
(4) no self-conflict, i.e. there is no $x \in B \cup E$ such that $x \#_{ON} x$;
(5) the set cut_0 of $\leqslant_{ON}$-minimal nodes is contained in B and finite.

In occurrence nets, the nodes of E are called *events*, and the elements of B are denoted *conditions*. Occurrence nets constitute particular cases of *prime event structures (PES)* in the sense of Winskel et al Nielsen et al. (1981); .

Definition 4. A *prime event structure (over alphabet $\mathbb{A}$)* is a tuple $\mathcal{E} = (E, \leqslant, \#)$, where E is a *set of events*,

(1) $\leqslant \subseteq E \times E$ is a partial order satisfying the property of *finite causes*, i.e. for all $e \in E$, $|[e]| < \infty$, and
(2) $\# \subseteq E \times E$ an irreflexive symmetric *conflict* relation satisfying the property of *conflict heredity*, i.e.

$$\forall\ e, e', e'' \in E :\ e \# e'\ \wedge e' \leqslant e'' \Rightarrow e \# e'', \quad (1)$$

Prefixes and Configurations. Restricting $\leqslant$ and $\#$ to the event set E, "forgetting" conditions of ON, yields an event structure. A *prefix* of $\mathcal{E}$ is any downward closed subset $V \subseteq E$, i.e. such that for every $e \in V$, $[e] \subseteq V$. Prefixes of $\mathcal{E}$ induce, in the obvious way, sub-event structures of $\mathcal{E}$ in the sense of the above definition. Denote the set of $\mathcal{E}$'s prefixes as $\mathcal{V}(\mathcal{E})$. Prefix $\mathbf{c} \in \mathcal{V}(\mathcal{E})$ is a *configuration* if and only if it is conflict-free, i.e. if $e \in \mathbf{c}$ and $e \# e'$ imply $e' \notin \mathbf{c}$. Denote as $\mathcal{C}(\mathcal{E})$ the set of $\mathcal{E}$'s configurations, and as $\mathcal{C}^{\mathbf{fin}}(\mathcal{E}) \subseteq \mathcal{C}(\mathcal{E})$ the set of all *finite* configurations. Call any $\subseteq$-maximal element of $\mathcal{C}(\mathcal{E})$ a *run* of $\mathcal{E}$; denote the set of $\mathcal{E}$'s runs as $\Omega(\mathcal{E})$, or simply Ω if no confusion can arise.

Definition 5. If $N_1 = (P_1, T_1, F_1)$ and $N_2 = (P_2, T_2, F_2)$ are nets, a *homomorphism* is a mapping $h : P_1 \cup T_1 \to P_2 \cup T_2$ such that *(i)* $h(P_1) \subseteq P_2$ and *(ii)* for every $t_1 \in T_1$, the restriction to ${}^\bullet t_1$ is a bijection between the set ${}^\bullet t_1$ in N_1 and the ${}^\bullet h(t_1)$ in N_2, and similarly for $t_1{}^\bullet$ and $(h(t_1))^\bullet$. A *branching process* of safe Petri net $\mathcal{N} = (N, M_0)$ is a pair $\beta = (ON, \pi)$, where $ON = (B, E, G)$ is an occurrence net, and π is a homomorphism from ON to N such that:

(1) The restriction of π to cut_0 is a bijection from cut_0 to M_0, and
(2) for every $e_1, e_2 \in E$, if ${}^\bullet e_1 = {}^\bullet e_2$ and $\beta(e_1) = \beta(e_2)$ then $e_1 = e_2$.

The unique (up to isomorphism) maximal branching process $\beta_{\mathbf{Unf}} = (ON_{\mathbf{Unf}}, \pi_{\mathbf{Unf}})$ of $\mathcal{N}$ is called the *unfolding* of $\mathcal{N}$; see Esparza and Vogler (2002) for a canonical algorithm to compute the unfolding of $\mathcal{N}$. We will assume that all transitions $t \in T$ have at least one output place, i.e. $t^\bullet$ is not empty. In this case, every finite configuration $\mathbf{c}$ of $ON_{\mathbf{Unf}}$ spans a conflict free subnet $\mathbf{c}_{\mathbf{Unf}} = (E_{\mathbf{c}}, B_{\mathbf{c}}, G_{|(E_{\mathbf{c}} \times B_{\mathbf{c}}) \cup (B_{\mathbf{c}} \times E_{\mathbf{c}})})$ of $ON_{\mathbf{Unf}}$ by setting

$$B_{\mathbf{c}} \triangleq \bigcup_{e \in E} ({}^\bullet e \cup e^\bullet).$$

The following results (see e.g. Esparza and Vogler (2002)) justify the use of unfoldings: The set $cut(\mathbf{c})$ of $\leqslant$-maximal nodes of $\mathbf{c}_{\mathbf{Unf}}$ is contained in $B_{\mathbf{c}}$. Moreover, $cut(\mathbf{c})$ is a *co-set*, that is, for all distinct conditions $b, b' \in cut(\mathbf{c})$, $b\ \mathbf{co}\ b'$ holds; and $cut(\mathbf{c})$ is $\subseteq -$maximal with this property, and such sets in occurrence nets are called *cuts*. By setting, for any cut cut and place p,

$$M_{cut}(p) \triangleq |\{b \in cut :\ \pi(b) = p\}|,$$

we obtain a marking of $\mathcal{N}$. Now, for $cut(\mathbf{c})$ as above, $M_{\mathbf{c}} \triangleq M_{cut(\mathbf{c})}$ is a reachable marking of $\mathcal{N}$, more precisely the marking that $\mathcal{N}$ is in after executing firable transitions in a sequence compatible with $\mathbf{c}$. Conversely, for every reachable marking M of $\mathcal{N}$ there exists (at least) one configuration $\mathbf{c}$ in $ON_{\mathbf{Unf}}$ such that $M_{\mathbf{c}} = M$.

Progressive configurations. For any finite configuration $\mathbf{c}$ and event $e \in E \backslash \mathbf{c}$ such that $\mathbf{c} \cup \{e\}$ is a configuration (in particular, there is no conflict between e and any event in $\mathbf{c}$ and all predecessors of e are contained in $\mathbf{c}$), we have that $M_{\mathbf{c}} \xrightarrow{\pi(e)}$. We therefore denote this situation by $\mathbf{c} \overset{e}{\hookrightarrow}$. Now, let the *height* of an event e be the longest $<$-chain of events leading to and including e:

(1) $\mathcal{H}(\emptyset) \triangleq 0$,
(2) $\mathcal{H}(e) \triangleq 1 + \max\{\mathcal{H}(e') :\ e' \in \langle e \rangle)\}$,

and for any configuration $\mathbf{c}$, let

$$\mathcal{H}(\mathbf{c}) \triangleq \sup_{e \in \mathbf{c}} \{\mathcal{H}(e)\}.$$

Then configuration $\mathbf{c}$ is called *progressive* iff for every $e \in E \backslash \mathbf{c}$ such that $\mathbf{c} \overset{e}{\hookrightarrow}$, one has $\mathcal{H}(\mathbf{c}) < \mathcal{H}(e)$. Denote by $\mathcal{C}_{\mathbf{prog}}$ the set of progressive configurations.

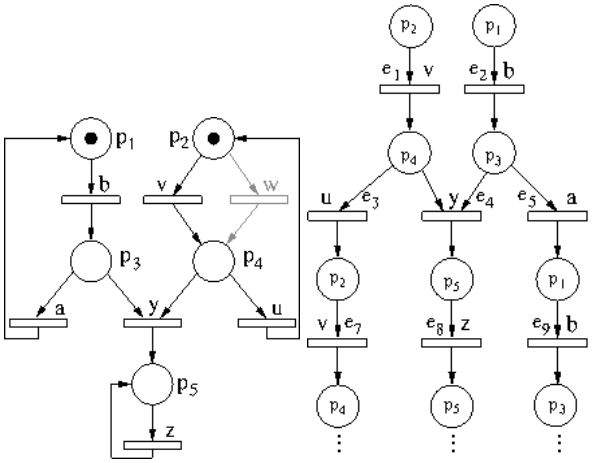

Fig. 1. Left: a Petri net $\mathcal{N}$ with two variants $\mathcal{N}_*$ and $\mathcal{N}_w$; as indicated by the grey lines, both nets are equal up to the presence of transition w in $\mathcal{N}_w$ and its absence from $\mathcal{N}_*$. On the right, a prefix of $\mathcal{N}_*$'s unfolding.

Complete Prefixes. Unfoldings of safe Petri nets are infinite in general. However, since the space of reachable markings is finite, all states and all patterns of behaviour can be observed on a bounded prefix of the unfolding. The shape and size of such a *complete prefix* varies depending on the information one wishes to extract, and on the method used to truncate the unfolding. Following Khomenko et al. (2003), we define:

Definition 6. A *cutting context* for $\mathbf{Unf}_\mathcal{N}$ is a triple $\Theta = (\sim, \prec, (\mathbf{c}_e)_{e\in E})$, where

(1) $\sim$ is an equivalence relation on $\mathcal{C}^{\mathbf{fin}}$,
(2) $\prec$ is a strict well-founded partial order on $\mathcal{C}^{\mathbf{fin}}$ such that $\mathbf{c} \subseteq \mathbf{c}'$ implies $\mathbf{c} \prec \mathbf{c}'$; $\prec$ is then called an *adequate* order;
(3) $\sim$ and $\prec$ are preserved by finite extensions; i.e. for every $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_1' \in \mathcal{C}^{\mathbf{fin}}$ such that *(i)* $\mathbf{c}_1 \sim \mathbf{c}_2$ and *(ii)* $\mathbf{c}_1 \subseteq \mathbf{c}_1'$, there exists $\mathbf{c}_2' \in \mathcal{C}^{\mathbf{fin}}$ with $\mathbf{c}_2 \subseteq \mathbf{c}_2'$ such that
   - $\mathbf{c}_1' \sim \mathbf{c}_2'$, and
   - $\mathbf{c}_1 \prec \mathbf{c}_2$ implies that $\mathbf{c}_1' \prec \mathbf{c}_2'$,
(4) and $(\mathcal{C}_e)_{e\in E}$ is a family of subsets $\mathcal{C}_e \subseteq \mathcal{C}^{\mathbf{fin}}$.

One defines recursively sets $\mathbf{coff}^\Theta$ and $\mathbf{fsb}^\Theta$, respectively of *cut-off* and *feasible* events, by :

(1) $e \in \mathbf{fsb}^\Theta$ iff $\langle e\rangle \cap \mathbf{coff}^\Theta = \emptyset$;
(2) e is a *static cut-off event* iff *(i)* e is feasible and *(ii)* there is a *corresponding* configuration $\mathbf{c} = \mathbf{c}(e) \in (\mathbf{c}_e)_{e\in E}$ such that
   - $\mathbf{c} \subseteq (\mathbf{fsb}^\Theta \backslash \mathbf{coff}^\Theta)$,
   - $\mathbf{c} \sim [e]$ and $\mathbf{c} \prec [e]$.

The branching process of $\mathcal{N}$ obtained by restricting to the events in $\mathbf{fsb}^\Theta$ is called the Θ-*canonical prefix* of $\mathbf{Unf}_\mathcal{N}$.

As shown in Khomenko et al. (2003), the canonical prefix is finite whenever there is no infinite $<$-chain of feasible events.

3. OBSERVABILITY AND DIAGNOSABILITY

Let $\mathcal{N} = (P, T, F, M_0)$ a safe Petri net, $\lambda : T \to \mathbb{A}$ a labeling mapping into an alphabet $\mathbb{A}$ that contains the empty symbol ε, $\mathbf{Unf}_\mathcal{N} = (B, E, G, cut_0)$ its unfolding net, with labeling morphism $\alpha : E \to T$ given by the unfolding morphism. Denote as $\mathcal{U} \triangleq \lambda^{-1}(\{\varepsilon\})$ the set of *unobservable* transitions, and as $\mathcal{O} \triangleq T\backslash\mathcal{U}$ the set of *observable* transitions; accordingly, let $E_\mathcal{U} \triangleq \alpha^{-1}(\mathcal{U})$ and $E_\mathcal{O} \triangleq \alpha^{-1}(\mathcal{O})$ be the set of unobservable and observable events of $\mathbf{Unf}_\mathcal{N}$, respectively.

Due to partial observation, different configurations can be equivalent in terms of the *observable* behaviour. The following notations formalize this fact.

Definition 7. Write $\mathbf{c} \sim_\mathbb{O} \mathbf{c}'$ iff for the restrictions of $\leqslant$ to $\mathbf{c}_\mathbb{O} \triangleq \mathbf{c} \cap E_\mathbb{O}$ and $\mathbf{c}'_\mathbb{O} \triangleq \mathbf{c}' \cap E_\mathbb{O}$, there exists a label-preserving order isomorphism $I : \mathbf{c}_\mathbb{O} \to \mathbf{c}'_\mathbb{O}$. If $\mathbf{c}_1 \subseteq \mathbf{c}_2$ and $\mathbf{c}_2 \sim_\mathbb{O} \mathbf{c}_3$, then write $\mathbf{c}_1 \subseteq_\mathbb{O} \mathbf{c}_3$.

Definition 8. Let $\sigma = t_1 t_2 \ldots \in T^\omega \triangleq T^* \cup T^\infty$ be a transition sequence that is enabled in M_0, i.e. assume there exist reachable markings $M_1, M_2, \ldots$ of $\mathcal{N}$ such that

$$M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \xrightarrow{t_3} \ldots$$

Then σ is *weakly fair* iff for any $t \in T$ and $i \in \mathbb{N}$ for which $M_i \xrightarrow{t}$, there exists $j > i$ such that $M \not\xrightarrow{t}$.

In other words, weakly fair executions are such that no transition remains enabled "forever": after any transition t's enabling on σ, t must eventually become disabled, either by its own firing, or by the firing of a conflicting transition. In the unfolding of $\mathcal{N}$, a fair run corresponds to a set of events ω such that for any event e, either $e \in \omega$, or there exists $e' \in \omega$ such that $e \# e'$. Of course, this is equivalent to ω being a maximal configuration, i.e. $\omega \in \Omega$. Let $\phi \in \mathcal{U}$ be a fault transition, and let $E_\phi \triangleq \alpha^{-1}(\phi)$. With these preparations, we are ready to define:

Definition 9. We say that $\mathcal{N}$ is **weakly observable w.r.t.** λ iff for every $\omega \in \Omega(\mathcal{N})$, $|\omega \cap E_\mathcal{O}| = \infty$. A weakly observable (w.r.t. λ) $\mathcal{N}$ is **weakly diagnosable** w.r.t. λ and ϕ iff for every faulty run $\omega_\phi \in \Omega(\mathcal{N})$, it holds that any $\omega \in \Omega(\mathcal{N})$ such that $\omega \sim_\mathcal{O} \omega_\phi$ satisfies $\omega \cap e_\phi \neq \emptyset$.

Example. In the net $\mathcal{N}_*$ in Figure 1 on the left , let the fault be $\phi = v$, and assume that a is the only observable transition. Then:

a) In sequential semantics, the run which consists only of occurrences of u and v is infinite but produces no observation; $\mathcal{N}_*$ is therefore not (strongly) observable in the classical sense. Moreover, $\mathcal{N}_*$ is not (strongly) diagnosable, since all runs *without* an occurrence of y are observationally indiscernable from the run ω' formed only by occurrences of a and b; there exist thus observationally equivalent runs some of which are faulty, and some healthy.

b) However, with the same assumptions, $\mathcal{N}_*$ is both *weakly observable* and *weakly diagnosable*. In fact, every run ω is fault-definite since v must have occured.

The next section will make these intuitions more precise. Here, let us make one further observation in the context of the example. In fact, in decentralized systems with weak synchronization between subsystems, faults may elude diagnosis under the interleaved viewpoint, while being *weakly* captured under partial order semantics. In the example, consider now b the fault event, instead of v, and let still a be observable. Then, the new system is neither

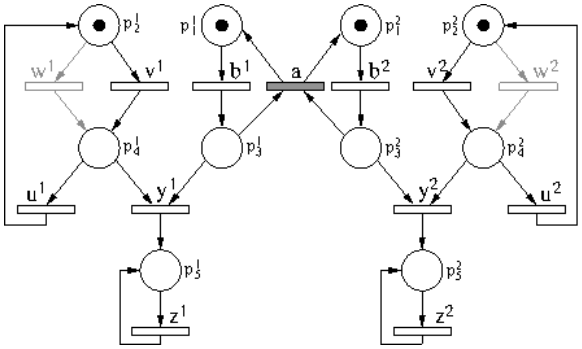

Fig. 2. The verifier net $\mathcal{V}$ of the running example synchronized on the observable transition a (highlighted). The superscript is used to distinguish nodes belonging to $\mathcal{N}_1$ and $\mathcal{N}_2$, respectively.

classically observable nor classically diagnosable. However, removing the loop $u-v$ from the system leaves a classically diagnosable system. In other words, it is the presence of the second loop, running in parallel and without influence on the fault occurrence, that blocks diagnosis of the fault. [1] Thus, the partial order approach actually *increases* precision for partial observation of highly concurrent systems.

4. VERIFICATION OF WEAK DIAGNOSABILITY

The Verifier Net. For the practical verification, we propose use an extension of the unfolding-based verifier method developped in Madalinski et al. (2010). The basic idea is to synchronize, via the observable labels, two copies of the supervised net $\mathcal{N}$ and to check for the existence of executions of the product net in which the projection to the first component is faulty while that to the second component is healthy. Formally, with the setup of Definition 7, let $\mathcal{V} \triangleq \mathcal{N}_1 \times \mathcal{N}_2$ be the α-synchronized product of two isomorphic copies $\mathcal{N}_1$ and $\mathcal{N}_2$ of $\mathcal{N}$, i.e. $\mathcal{N}_i = (P_i, T_i, F_i, M_0^i)$. That is, with

(1) $P_\mathcal{V} = p_1 \uplus p_2$,
(2) for $i \in \{1,2\}$, $T_i^\varepsilon \triangleq \{t \in T_i \mid \alpha(t) = \varepsilon\}$,
(3) $T_{12} \triangleq t\{t \in T_1 \mid \alpha(t) \neq \varepsilon\}$,
(4) $F_\varepsilon \triangleq \bigcup_{i=1}^2 (F_i \cap P_i \times T_i^\varepsilon) \cup \bigcup_{i=1}^2 (F_i \cap T_i^\varepsilon \times P_i)$,
(5) $F_{12} \triangleq \bigcup_{i=1}^2 (F_i \cap P_i \times T_{12}) \cup \bigcup_{i=1}^2 (F_i \cap T_{12} \times P_i)$,
(6) $T_\mathcal{V} \triangleq T_1^\varepsilon \uplus T_2^\varepsilon \uplus T_{12}$ and $F_\mathcal{V} \triangleq F_\varepsilon \uplus F_{12}$,
(7) $M_0 \triangleq M_0^1 \uplus M_0^2$,

$\mathcal{V}$ is the Petri net $\mathcal{N}_\mathcal{V} = (p_\mathcal{V}, T_\mathcal{V}, F_\mathcal{V}, M_0)$, with the labeling $\alpha : T_\mathcal{V} \to \mathbb{A}$ inherited from $\mathcal{N}$. The verifier of the running example is depicted in Figure 2.

Lemma 1. The configurations of $\mathcal{V}$ are given by pairs $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$ of configurations of $\mathcal{N}_1, \mathcal{N}_2$, respectively, where

(1) $\mathbf{c}_i$ is the projection of $\mathbf{c}$ to the occurrences and conditions for $\mathcal{N}_i$, and
(2) there exists a partial mapping $\psi : \mathbf{c} \to \mathbf{c}'$ such that
   - setting $\mathbf{c}_\mathcal{O} \triangleq E_\mathcal{O} \cap \mathbf{c}$ and $\mathbf{c}_\mathcal{U} \triangleq E_\mathcal{U} \cap \mathbf{c}$ (and analogously for $\mathbf{c}'$), $\psi(e)$ is defined for all $e \in \mathbf{c}_\mathcal{O}$, and *undefined* for all $e \in \mathbf{c}_\mathcal{U}$;
   - the restriction $\psi_{|\mathbf{c}_\mathcal{O}}$ defines a bijection $\psi_{|\mathbf{c}_\mathcal{O}} : \mathbf{c}_\mathcal{O} \to \mathbf{c}'_\mathcal{O}$,
   - for all $e_1, e_2 \in \mathbf{c}_\mathcal{O}$ such that $e_1 \neq e_2$, if $e_1 < e_2$ then $\neg(\psi(e_2) < \psi(e_1))$.

[1] Thanks to A. Guia who made us discover this aspect by a remark in a workshop discussion with the third author.

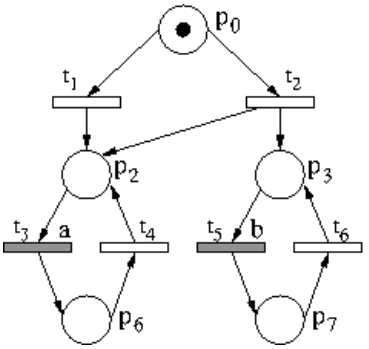

Fig. 3. A safe Petri net having two weakly fair runs ω_1, ω_2 such that the observable image $\lambda(\omega_1)$ of ω_1 is a proper prefix of $\lambda(\omega_2)$: ω_1 has only one interleaving $\sigma_1 = t_1 t_3 t_4 t_3 t_4 t_3 t_4 \ldots$, while ω_2 is formed by one occurrence of t_2 and infinitely many occurrences of t_5 and t_6.

Proof: The decomposition of $\mathbf{c}$ follows from the synchronized product and the construction of $\mathbf{Unf}_\mathcal{V}$; since $\mathbf{c}$ is free of cycles, it follows that $e_1 < e_2$ implies. □

For the construction of a verifier, we follow Madalinski et al. (2010) in using a fault indicator variable whose value is one for all configurations on which a fault occurs, and 0 otherwise. For this, let $\Phi : E \to \{0,1\}$ be such that $\Phi(e) = 1$ if $e \in e_\Phi$, and 0 otherwise; then, define recursively $\nu : E \to \{0,1\}$ by $\nu(cut_0) = 0$ and

$$\nu(e) \triangleq \max\left[\Phi(e), \max_{e'<e} \nu(e')\right]$$

Mapping ν extends naturally to a mapping $\nu : \mathcal{C} \to \{0,1\}$ by setting $\nu(\mathbf{c}) \triangleq \sup_{e \in \mathbf{c}} \nu(e)$. Moreover, recall that every configuration $\mathbf{c}$ of $\mathbf{Unf}_\mathcal{V}$ is given as a pair $(\mathbf{c}_1, \mathbf{c}_2)$ of configurations of $\mathbf{Unf}_\mathcal{N}$; we therefore have a 2-vector valued mapping $\boldsymbol{\nu} : \mathcal{C}(\mathcal{V}) \to \{0,1\}^2$ given by $\boldsymbol{\nu}(\mathbf{c}) \triangleq (\nu(\mathbf{c}_1), \nu(\mathbf{c}_2))$. *The canonical prefix* $\mathbf{V}_\mathcal{V}$. For the practical verification of weak diagnosability, we need to adapt our choice of cutting context to obtain a sufficient finite prefix of the verifier. The crucial point is the choice of the collection $\mathcal{C}_e \subseteq \mathcal{C}^{\mathbf{fin}}$ for each event. Denote by

$$\overline{\mathcal{C}^e_{\mathbf{prog}}} \triangleq \{\mathbf{c} \in \mathcal{C}_{\mathbf{prog}} \mid e \in \mathbf{c}\}$$

the set of *progressive* configurations containing event e, i.e. the progressive extensions of $[e]$, and by

$$\mathcal{C}^e_{\mathbf{prog}} \triangleq \left\{\mathbf{c} \in \overline{\mathcal{C}^e_{\mathbf{prog}}} \mid \forall \mathbf{c}' \in \overline{\mathcal{C}^e_{\mathbf{prog}}} :\ \mathbf{c}' \subseteq \mathbf{c} \Rightarrow \mathbf{c}' = \mathbf{c}\right\}$$

the set of progressive $[e]$-extensions that are minimal with this property. First, we note that for $e \in E$, $\mathcal{C}^e_{\mathbf{prog}} \neq \emptyset$; in fact, one obtains all configurations of $\mathcal{C}^e_{\mathbf{prog}}$ by the following non-deterministic algorithm:

- Set $N \triangleq \mathcal{H}(e)$ and $\mathbf{c}_0 \triangleq [e]$.
- For $n \geqslant 1$, set $E_n \triangleq \{e \in E \backslash \mathbf{c}_{n-1} \mid \mathcal{H}(e) \leqslant n\}$
- Choose $e \in E_n$ such that $\langle e \rangle \subseteq \mathbf{c}_{n-1}$ (i.e. $\mathbf{c}_{n-1} \overset{e}{\hookrightarrow}$ - and set $\mathbf{c}_n \triangleq \mathbf{c}_{n-1} \cup \{e\}$.
- Repeat until $E_n = \emptyset$.

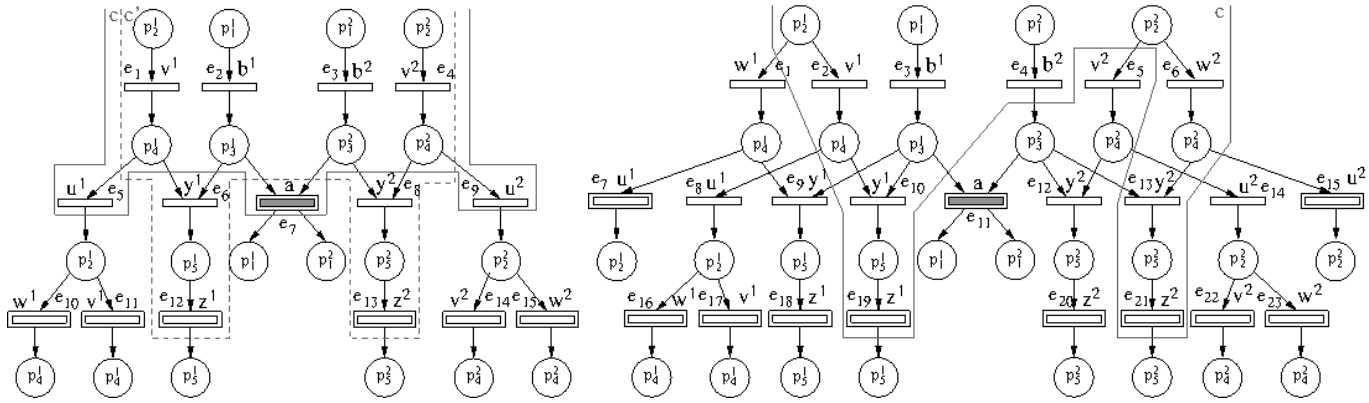

Fig. 4. The complete prefix of the running examples $\mathcal{N}_*$ (left) and $\mathcal{N}_w$ (right) from Figure 1. The cut-off events are represented with double boxes. Both are weakly observable; the left prefix - for $\mathcal{N}_*$- shows weak diagnosability, whereas the right prefix - for $\mathcal{N}_w$ - does not.

Then, the definition of the cutting context for verification of weak diagnosability reads as follows:

(1) $\mathbf{c} \sim \mathbf{c}'$ iff *(i)* $M_{\mathbf{c}} = M_{\mathbf{c}'}$ and *(ii)* $\phi \in \pi(\mathbf{c})$ iff $\phi \in \pi(\mathbf{c}')$;
(2) $\mathbf{c} \prec \mathbf{c}'$ iff $\mathbf{c} \subseteq \mathbf{c}'$ and
(3) $\mathcal{C}_e \triangleq \mathcal{C}^e_{\mathbf{prog}}$.

One checks that the relations $\sim$ and $\prec$ thus defined satisfy the conditions of Definition 6. Moreover:

Lemma 2. There is no infinite $<$-chain of feasible events.

Proof: Assume there exist such $e_1 < e_2 < \ldots$. Since $\sim$ has only finitely many distinct classes in $\mathcal{C}$ (a fact that follows from 1-safeness), there must exist $i < j$ such that $[e_i] \sim [e_j]$. Therefore one finds $\mathbf{c}_i \in \mathcal{C}_{e_i}$ and $\mathbf{c}_j \in \mathcal{C}_{e_j}$ such that $\mathbf{c}_i \sim \mathbf{c}_j$ and $\mathbf{c}_i \subseteq \mathbf{c}_j$, which contradicts the feasibility of e_j. □

Therefore, by Khomenko et al. (2003), we obtain a complete canonical finite prefix $\mathbf{V}_{\mathcal{N}}$ of $\mathcal{N}$ for any safe Petri net $\mathcal{N}$. Specializing to the verifier net $\mathcal{V}$ defined above, denote $\mathbf{V}_{\mathcal{V}}$ by $\mathbf{V}$ for simplicity.

We first observe:

Theorem 1. With the above notation, let $\mathbf{Unf}_{\mathcal{V}}$ be the unfolding of $\mathcal{V}$. Then $\mathcal{N}$ is weakly diagnosable iff for every progressive configuration $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$ of $\mathcal{V}$ such that both $\mathbf{c}_1$ and $\mathbf{c}_2$ are progressive for $\mathcal{N}$, either both $\mathbf{c}_1$ and $\mathbf{c}_2$ are faulty, or both are healthy, that is: $\boldsymbol{\nu}(\mathbf{c}) \in \{(0,0),(1,1)\}$.

Proof: Follows directly from the definitions. □

We have :

Theorem 2. Assume $\mathcal{N}$ is weakly observable. If for every $e \in \mathbf{coff}(\mathcal{V})$, one has $\boldsymbol{\nu}([e]) \in \{(0,0),(1,1)\}$, then $\mathcal{N}$ is weakly diagnosable.

Proof: Assume that all cut-off events e satisfy $\boldsymbol{\nu}([e]) \in \{(0,0),(1,1)\}$, but that there exists $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$ such that, w.l.o.g., $\boldsymbol{\nu}(\mathbf{c}) = (1,0)$ and $\mathbf{c}_1$ is progressive. Choose $\mathbf{c}$ $\prec$-minimal with this property, and let ϕ be the $<$-minimal occurrence of a fault event e_ϕ in $\mathbf{c}$. By construction of $\mathbf{V}$ and Theorem 1, there exists a corresponding event e'_ϕ in $\mathbf{V}$; if e'_ϕ is a cut-off event, we have a contradiction, since $\boldsymbol{\nu}(e'_\phi) = (1,0)$ by construction. Hence assume e'_ϕ is not cut-off; then there is $\mathbf{c}_{e'_\phi} \in \mathcal{C}_e$ such that $\mathbf{c}_{e'_\phi} \prec \mathbf{c}$, contradicting the $\prec$-minimal choice of $\mathbf{c}$, and we are done. □

The converse of Theorem 1 does *not hold*. In fact, note first that there are in general progressive configurations of $\mathcal{V}$ that do *not* project to progressve configurations of the components $\mathcal{N}_1$ and $\mathcal{N}_2$. This is the case for the net $\mathcal{N}$ shown in figure 3: in $\mathcal{N}$, with a and b the only observable transitions, take the configuration $\mathbf{c}_1$ obtained by firing t_1, t_3, t_4 exactly once, and let $\overline{\mathbf{c}}_2$ be the configuration obtained by one firing each of t_2, t_3, t_4, t_5, t_6. Now, the verifier (which we do not draw here due to space limitations) $\mathcal{V}$ has a progressive configuration $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2)$ whose projection to $\mathcal{N}_1$ is $\mathbf{c}_1$ and whose $\mathcal{N}_2$-image is the configuration $\mathbf{c}_2$ obtained by firing t_2, t_3, t_4 exactly once. Clearly, $\mathbf{c}_2$ is a proper prefix of $\overline{\mathbf{c}}_2$, and $\mathbf{c}_2$ is *not* progressive. This means that, for a general net $\mathcal{N}$, the verifier $\mathcal{V}$'s verdict provides a semi-decision, which has to be complemented:

(1) If there is no ambiguous $\mathbf{c} \in \mathcal{C}_{\mathbf{prog}}$, then $\mathcal{N}$ is weakly diagnosable.
(2) If $\mathcal{V}$ exhibits an *ambiguity witness*, i.e.$\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2) \in \mathbf{c}_{\mathbf{prog}}(\mathcal{V})$ such that (w.l.o.g) $\mathbf{c}_1$ is a *faulty* configuration of $\mathcal{N}_1$ and $\mathbf{c}_2$ a *healthy* configuration of $\mathcal{N}_2$, it must be verified (this can be done on finite prefixes of $\mathcal{N}$ whose size is bounded by that of $\mathbf{V}(\mathcal{V})$) whether *both* $\mathbf{c}_1 \in \mathcal{C}_{\mathbf{prog}}(\mathcal{N}_1)$ *and* $\mathbf{c}_2 \in \mathcal{C}_{\mathbf{prog}}(\mathcal{N}_2)$ hold. If so , then $\mathcal{N}$ is *not* weakly diagnosable. Otherwise, if there is another witness from $\mathcal{V}$, inspect that witness; otherwise $\mathcal{N}$ *is* weakly diagnosable.

The net $\mathcal{N}_*$ of the running example (Figure 1 on the left) is weakly diagnosable. This can be analyzed on the verifier's unfolding prefix depicted in Figure 4 on the left. To avoid a lengthy enumeration, consider the following informal analysis: there exist maximal configurations of the verifier prefix with *(i)* one or *(ii)* zero occurrences of a. With one occurrence of a, we must have also, in every maximal configuration of $\mathbf{V}$, one occurrence each of v and v'. This is reflected by the cut-off $\mathbf{c} = \{e_1, e_2, e_3, e_4, e_5, e_7, e_9\}$ with $\boldsymbol{\nu}(\mathbf{c}) = (1,1)$. Finally, if a does not occur, then we must have occurrence of the highest (in the figure) instances of y and y', which is only possible if there is exactly one occurrence each of v and v'. This is illustrated by the cut-off $\mathbf{c}' = \{e_1, e_2, e_3, e_4, e_6, e_8, e_{12}, e_{13}\}$ with $\boldsymbol{\nu}(\mathbf{c}') = (1,1)$.

Consider $\mathcal{N}_w$ from Figure 4. On the right hand side, $\mathbf{c} = \{e_2, e_3, e_4, e_6, e_{10}, e_{13}, e_{19}, e_{21}\}$ with $\boldsymbol{\nu}(\mathbf{c}) = (1,0)$ witnesses a violation of weak diagnosability. In fact, in

the entire example one cannot distinguish between runs on which only w occurs - that is, *healthy* runs - from the faulty ones that contain occurrences of v; $\mathcal{N}_w$ is not weakly diagnosable.

5. CONCLUSION AND OUTLOOK

We have provided a cornerstone for partial order diagnosis for safe Petri nets, by showing how weak diagnosability can be *effectively* verified using a finite occurrence net. The main construction is that of a complete finite prefix of the unfolding of the *verifier net* obtained as the product of to copies of the system model $\mathcal{N}$, synchronized by fusing only *observable* transitions.

In Madalinski et al. (2010), the verifier construction with Petri net unfoldings, on which our approach builds up, had been developped in the context of verification of *strong* diagnosability. Moving to the problem of *weak* diagnosability required a subtle modification of the cutting context. In fact, unfoldings are most frequently exploited, and cut off, using *prime configurations* $[e]$ only; this was shown in Madalinski et al. (2010) to capture efficiently violations of (strong) diagnosability. However, these configurations are in general not progressive, and do not allow to detect faults in unsynchronized parts of the net, such as in $\mathcal{N}_*$ above. For analyzing nets that are not strongly diagnosable but might still allow weak diagnosis ("based on the observation, v is eventually inevitable"), like $\mathcal{N}_*$, the system of prime configurations is not adequate. The key to extending the verifier approach was therefore the adaptation of the cutting context, in the sense of Khomenko et al. (2003), so that the cut-off criteria could be based on a suitable collection of finite progressive configurations. Showing the validity of the adapted verifier approach for weak diagnosability is the main contribution of the paper.

The *efficiency* of the unfolding-based construction - which is PSCPACE-complete in general - hinges upon the size of the complete prefix , and thus upon the wise choice of cut-off context. Here, a very conservative adequate order was chosen, which orders configurations merely by inclusion; exploring more sophisticated ordering relations can be a source of important space reductions. Future work will explore different choices of such cutting contexts adapted to the weak diagnosability setting.

More generally, there is room to explore further improvements in the exploration and storage of $\mathbf{V}$. In fact, the prefixes proposed above tend to have greater width that those obtained with prime configuration-based cutting criteria. We will strive to identify efficient techniques for pruning away unnecessary branches at as early a stage as possible.

Another approach to partial observation in concurrent systems, which has been introduced in Haar (2007, 2009, 2010a,b), consists in looking for *inevitable* occurrences that are revealed by observation, regardless of the possible time for occurrence (which may be concurrent with the observation, with no synchronization). Knowledge of such relations in the system allows to raise alarms and start countermeasures as soon as the threat becomes apparent, without waiting for evidence of its actual occurrence.

Acknowledgments: This work was partly supported by the European Community's 7th Framework Programme under project DISC (*DI*stributed *S*upervisor *C*ontrol of large plants), Grant Agreement INFSO-ICT-224498, the Fondecyt project No. 11090257, and the *ARCUS* project *Ile de France/Inde*, conv. F-68-1309/R.

REFERENCES

Paolo Baldan, Thomas Chatain, Stefan Haar, and Barbara König. Unfolding-based diagnosis of systems with an evolving topology. *Information and Computation*, 208 (10):1169–1192, October 2010.

Albert Benveniste, Éric Fabre, Stefan Haar, and Claude Jard. Diagnosis of asynchronous discrete event systems: A net unfolding approach. *IEEE Transactions on Automatic Control*, 48(5):714–727, May 2003.

C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, Boston etc, 1999.

Stefan Haar. Unfold and cover: Qualitative diagnosability for Petri nets. In: *Proc. 46th IEEE CDC*, pp. 1886–1891, New Orleans, LA, USA, December 2007. IEEE Control System Society.

Stefan Haar. Qualitative diagnosability of labeled Petri nets revisited. In: *Proc. Joint 48th IEEE Conference on Decision and Control (CDC'09) and 28th Chinese Control Conference (CCC'09)*, pp. 1248–1253, Shanghai, China, December 2009. IEEE Control System Society.

Stefan Haar. Types of asynchronous diagnosability and the *reveals*-relation in occurrence nets. *IEEE Transactions on Automatic Control*, 55(10):2310–2320, October 2010.

Stefan Haar. What topology tells us about diagnosability in partial order semantics. In. *Proc. 10th Workshop on Discrete Event Systems (WODES'10)*, 2010.

Stefan Haar, Albert Benveniste, Éric Fabre, and Claude Jard. Partial order diagnosability of discrete event systems using Petri net unfoldings. In *Proc. 42nd CDC'03*, vol 4, pp. 3748–3753, Hawaii, USA, December 2003. IEEE Control System Society.

S. Römer J. Esparza and W. Vogler. An improvement of McMillan's unfolding algorithm. *Formal Methods in System Design* **20**(3):285–310, 2002.

V. Khomenko and M. Koutny and W. Vogler. Canonical Prefixes of Petri net unfoldings. *Acta Informatica* **40**(2):95–118, 2003.

Agnes Madalinski, Farid Nouioua, and Philippe Dague. Diagnosability verification with Petri net unfoldings. *KES Journal*, 14(2):49–55, 2010. Long version: RR No. 1516, UMR 8623, CNRS, UParis-Sud, March 2009.

M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures, and domains (I). *Theoretical Computer Science*, 13:85–108, 1981.

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

G. Winskel. Event structures. In *Advances in Petri nets*, LNCS **255**, pp. 325–392, 1987, DOI: 10.1007/3-540-17906-2. Springer Verlag.