Preprints of WODES 2010
August 30th – September 1st, 2010
Berlin, Germany

What Topology tells us about Diagnosability in partial order semantics

Stefan Haar *

* *INRIA and LSV (CNRS-ENS Cachan)*
61, avenue du Président Wilson
94235 CACHAN Cedex - France
(e-mail:haar@lsv.ens-cachan.fr,stefan.haar@inria.fr).

Abstract: From a partial observation of the behaviour of a labeled Discrete Event System, *fault Diagnosis* strives to determine whether or not a given "invisible" fault event has occurred. The *diagnosability problem* can be stated as follows: does the labeling allow for an outside observer to determine the occurrence of the fault, no later than a bounded number of events after that unobservable occurrence ? In concurrent systems, partial order semantics adds to the difficulty of the problem, but also provides a richer and more complex picture of observation and diagnosis. In particular, it is crucial to clarify the intuitive notion of "*time after fault occurrence*". To this end, we will use a unifying metric framework for event structures, providing a general topological description of diagnosability in both sequential and nonsequential semantics for Petri nets.

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, (behavioural) language $\mathcal{L} \subseteq T^*$, and a set of *observable transition labels* $\mathbb{O}$. The associated labeling map, let us call it $\eta : T \to \mathbb{O}$ in line with the formalism used below, 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 extending η into a homomorphism $T^* \to \mathbb{O}^*$. A classical definition of diagnosability is given in [(4)], following [(27)]; writing $s \sim_\eta s'$ iff $s, s' \in T^*$ are mapped to the same observable word in $\mathbb{O}^*$, we can state it as follows:

$\mathcal{L}$ is *non-diagnosable* iff there exist sequences $s_N, s_Y \in \mathcal{L}$ such that:

(1) s_Y is faulty, s_N is healthy, and $s_N \sim_\eta s_Y$;
(2) moreover, s_Y with the above is 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, it also makes sense *conceptually* to use models that reflect the local and distributed nature of the observed system, such as Petri nets or graph grammars. Putting these ideas together, we were led in [(8)] to carry over diagnosis to asynchronous models *and their non-interleaved semantics*; see also the discussion of the necessity for using partial order methods in [(7)]. This generalized methodology for fault diagnosis is based on the non-sequential executions of labeled Petri nets. We have provided a series of results [(14; 16; 17; 18)] on *partial order diagnosability* for Petri nets, in the spirit of the above definition. 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 wea diagnosability notions in [(14; 18)].

Bauer and Pinchinat [(2)] have given a topological view on diagnosability in terms of sequential languages. Their results are confirmed and generalized in the present work. In fact, the topological framework is obtained applying suitable metrics to event structures; for this, we generalize a standard metric construction to be found in [(3; 20)] and others, in such a way that progress and observation properties can be captured in the resulting topology. Event structures provide a unifying semantical model both for the sequential and non-sequential viewpoints. That is, both sequential languages as in [(4; 2)] AND the partial order semantics given in [(5; 25)] and used in [(9; 18)] associate event structures to a system; and the metric topology given here coincides, on the sequential semantics, with the Cantor topology used in [(2)]. The characterization of diagnosable systems from [(18)] is generalized in the present article to the topological level, and beyond the domain of *safe* Petri nets and finite state systems.

Structure of the paper: We begin in Section 2. with the basic definitions for (labeled) event structures, and provide the link with Petri nets in Section 3. The following Section

4. investigates partial observation and diagnosability, and develops the main contributions of this paper. We conclude in Section 5.

2. EVENT STRUCTURES

Preliminaries. Let A be a set. $A^* \triangleq \{a_1 \dots a_n \mid a_i \in A\}$ is the set of all finite words over A; the set of *infinite* words over A is denoted A^ω. Let $\mathbf{1}_A$ be the indicator function of A, i.e. $\mathbf{1}_A(x) = 1$ iff $x \in A$ and $\mathbf{1}_A(x) = 0$ for $x \notin A$. Let $f : A \to B$ be a partial function. Write $f(a)\downarrow$ if f is defined on $a \in A$, and $f(a)\uparrow$ otherwise. The *domain* of f is $dom(f) \triangleq \{a \in A \mid f(a)\downarrow\}$, and the *image* of f is

$$f(A) \triangleq \{b \in B \mid \exists\, a \in dom(f) :\ f(a)\downarrow\ \wedge f(a) = b\}.$$

Event Structures. We shall be using throughout this paper *prime event structures (PES)* following Winskel et al [(25; 30)], with particular attention to labeling.

Definition 1. A *(labeled) prime event structure (over alphabet $\mathbb{A}$)* is a tuple $\mathcal{E} = (E, \leqslant, \#, \lambda)$, where

(1) $E = supp(\mathcal{E})$ is the *support*, or *set of events* of $\mathcal{E}$,
(2) $\leqslant \subseteq E \times E$ is a partial order satisfying the property of *finite causes*, i.e.
$$\forall\ e \in E :\ |\{e' \in E \mid e' \leqslant e\}| < \infty, \tag{1}$$
(3) $\# \subseteq E \times E$ an irreflexive symmetric *conflict* relation satisfying the property of *conflict heredity*, i.e.
$$\forall\ e, e', e'' \in E :\ e\#e'\ \wedge e' \leqslant e'' \Rightarrow e \leqslant e'', \tag{2}$$
(4) $\lambda : E \to \mathbb{A}$ is a total mapping called the *labelling.*

Events $e, e' \in E$ are *concurrent*, written $e\ \mathbf{co}\ e'$, iff neither $e = e'$ nor $e \leqslant e'$ $e' \leqslant= e$ nor $e\#e'$ hold. If $\mathbf{co} = \perp$, i.e. if $\mathbf{co}$ is the empty relation, we call $\mathcal{E}$ *sequential*. An $\mathbb{A}$-labeled event structure is called *simple* [1] iff

$$e\ \mathbf{co}\ e' \Rightarrow \lambda(e) \neq \lambda(e'). \tag{3}$$

A simple labeled event structure will be called an *SES*. Let $\mathcal{E}_1 = (E_1, \leqslant_1, \#_1, \lambda_1)$ and $\mathcal{E}_2 = (E_2, \leqslant_2, \#_2, \lambda_2)$ be two $\mathbb{A}$-labeled event structures.

(1) If $E_1 \subseteq E_2$ and for all $e, e' \in E_1$,
$$e\#_1 e' \Leftrightarrow e\#_2 e' \text{ and } e \leqslant_1 e' \Leftrightarrow e \leqslant_2 e',$$
then $\mathcal{E}_1$ is a *sub-event structure* of $\mathcal{E}_2$.
(2) A partial mapping $f : E_1 \to E_2$ is called an $(\mathbb{A}-)$ *morphism* iff
- $e \leqslant_1 e' \Rightarrow f(e) \leqslant_2 f(e')$
- $dom\,(\lambda_1) \subseteq dom\,(f)$ and $dom\,(f) \subseteq dom\,(\lambda_2)$
- $\forall\ e \in E_1 :\ \lambda_1(e) = \lambda_2\,(f\,(e))$.

(3) $\mathcal{E}_1$ and $\mathcal{E}_2$ are *$(\mathbb{A}-)$isomorphic* iff there exist morphisms $f : E_1 \to E_2$ and $f^{-1} : E_2 \to E_1$ such that for all $e_1 \in dom(f)$ and all $e_2 \in dom(f^{-1})$,
$$f^-\,(f\,(e_1)) = e_1 \text{ and } f\left(f^{-1}\,(e_2)\right) = e_2.$$

The *set of causes* or *prime configuration* of $e \in E$ is $[e] \triangleq \{e' \mid e' \leqslant e\}$. A *prefix* of $\mathcal{E}$ is any downward closed subset $B \subseteq E$, i.e. such that for every $e \in B$, $[e] \subseteq B$. Denote the set of $\mathcal{E}$'s prefixes as $\mathcal{B}(\mathcal{E})$. Prefix $\mathbf{c}$ is a *configuration* if and only if it is conflict-free, i.e. $e \in \mathbf{c}$ and $e\#e'$ imply $e' \notin \mathbf{c}$. Denote as $\mathcal{C}(\mathcal{E})$ the set of $\mathcal{E}$'s configurations. 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.

[1] one might call it *safe* or *auto-concurrency free*

Note that all prefixes of $\mathcal{E}$, and in particular all its configurations, constitute sub-event-structures of $\mathcal{E}$; we will denote these structures with the same symbols as the corresponding sets. For $\mathbf{c} \in \mathcal{C}$ and $S \subseteq \mathcal{C}$, let

$$\mathcal{C}_{\mathbf{c}} \triangleq \{\tilde{\mathbf{c}} \in \mathcal{C} \mid \mathbf{c} \subseteq \tilde{\mathbf{c}}\}\ ,\ \Omega_{\mathbf{c}} \triangleq \{\omega \in \Omega \mid \mathbf{c} \subseteq \omega\}$$
$$\text{and} \qquad \Omega_S \triangleq \bigcup_{\mathbf{c} \in S} \Omega_{\mathbf{c}}.$$

Further, for any $\mathbf{c} \in \mathcal{C}(\mathcal{E})$, denote as

$$\mathcal{E}_{\mathbf{c}} = (E_{\mathbf{c}}, \leqslant_{|E_{\mathbf{c}}}, \#_{|E_{\mathbf{c}}}, \lambda_{|E_{\mathbf{c}}}),$$
$$\text{where } E_{\mathbf{c}} \triangleq \{e \in E \backslash \mathbf{c} \mid \forall\ e' \in \mathbf{c} :\ \neg(e\ \#\ e')\},$$

the *shift* of $\mathcal{E}$ by $\mathbf{c}$. If $\mathbf{c}' \in \mathcal{C}(\mathcal{E}_{\mathbf{c}})$, then $\mathbf{c} \circ \mathbf{c}'$ is the unique configuration of $\mathcal{E}$ such that *(i)* $\mathbf{c}$ is a prefix of $\mathbf{c} \circ \mathbf{c}'$, and *(ii)* $\mathbf{c} \circ \mathbf{c}' \cap E_{\mathbf{c}} = \mathbf{c}'$. For every $\mathbf{c}' \in \mathcal{C}(\mathcal{E}_{\mathbf{c}})$, we observe that $\mathbf{c}'' \triangleq \mathbf{c} \cup \mathbf{c}' \in \mathbf{c}(\mathcal{E})$; write in this case $\mathbf{c}'' = \mathbf{c} \circ \mathbf{c}'$, and say that $\mathbf{c}''$ is obtained by *appending* $\mathbf{c}'$ to $\mathbf{c}$.

Prefix relations. Let $[\mathcal{E}]_{\mathbb{A}}$ be the $\mathbb{A}$-isomorphism class of $\mathbb{A}$-labeled PES $\mathcal{E}$. Denote

$$\mathcal{B}_{\mathbb{A}}(\mathcal{E}) \triangleq \{[B]_{\mathbb{A}} \mid B \in \mathcal{B}(\mathcal{E})\} \tag{4}$$
$$\mathcal{C}_{\mathbb{A}}(\mathcal{E}) \triangleq \{[\mathbf{c}]_{\mathbb{A}} \mid \mathbf{c} \in \mathcal{C}(\mathcal{E})\} \tag{5}$$

Write $B_1 \sqsubseteq B_2$ iff B_1 is $\mathbb{A}$-isomorphic to a prefix of B_2; $\sqsubseteq$ lifts to a binary relation on $\mathcal{B}_{\mathbb{A}}(\mathcal{E})$, which we will denote by the same symbol $\sqsubseteq$. For $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}(\mathcal{E})$, let $[\mathbf{c}_1] \sqcap [\mathbf{c}_2] \triangleq [\mathbf{c}_3]$, where $\mathbf{c}_3$ is the $\subseteq$-maximal prefix of $\mathbf{c}_1$ such that $\mathbf{c}_2$ has a prefix $\mathbf{c}_3'$ that is $\mathbb{A}$-isomorphic to $\mathbf{c}_3$.

Metrics. The sets $\mathcal{C}(\mathcal{E})$ and $\Omega(\mathcal{E})$ can be equipped with Lawson or Scott topologies, or with natural *metrics*; we will follow and generalize the latter approach, similar to metrizations of traces as studied in [(19)]. Our pseudometrics allow to capture in particular partial observation and fault equivalence. Our principal tool are **μ-Heights**: Let $\mu : \mathbb{A} \to \mathbb{R}_0^+$ be any total mapping; we shall refer to μ as a *weight* function. As a particular case, consider $\mu(e) \equiv \mathbf{1}_E$: we will refer this as the *counting weight*. The following construction yields pseudometrics that are equivalent (in topological terms) to the prefix metric [(20)] and the Foata normal form metric [(3)], see [(19)], when the counting weight is chosen; other choices of weights allow to generalize to observation and fault equivalence.

The *μ-induced $*$-height* $\mathcal{H}^*_\mu(B)$ of a *prefix* is defined recursively by setting, for $\emptyset$ representing the empty preset,

$$\mathcal{H}^*_\mu(\emptyset) \triangleq 0 \tag{6}$$
$$\mathcal{H}^*_\mu([e]) \triangleq \mathcal{H}^*_\mu([e] \setminus \{e\}) + \mu(e) \tag{7}$$
$$\mathcal{H}^*_\mu(B) \triangleq \sup_{e \in B}(\mathcal{H}^*_\mu([e])). \tag{8}$$

Now, for $\tau \in [0, \infty)$ let $\mathcal{U}^\mu_\tau$ be the *τ-prefix under μ*, i.e.

$$\mathcal{U}^\mu_\tau \triangleq \bigcup \left\{B \in \mathcal{B}(\mathcal{E}) \mid \mathcal{H}^*_\mu(B) \leqslant \tau\right\}, \tag{9}$$

and let $\mathcal{E}^\mu_\tau$ be the prime event structure that $\mathcal{E}$ induces on $\mathcal{U}^\mu_\tau$. Then define $\mathcal{H}_\mu(\mathbf{c})$ for all $\mathbf{c} \in \mathcal{C}(\mathcal{E})$ as

$$\mathcal{H}_\mu(\mathbf{c}) \triangleq \sup\{\tau \mid \mathbf{c} \in \Omega(\mathcal{E}^\mu_\tau)\}. \tag{10}$$

Note that $\mathcal{H}_\mu(\bullet)$ is invariant under $\mathbb{A}$-isomorphism. Thus, let $\Psi_\mu(\bullet) : \mathcal{C}(\mathcal{E}) \to [0,1]$ and the *μ-pseudometric* $\mathbf{d}_\mu(\bullet,\bullet)$ be given by

$$\Psi_\mu(\mathbf{c}) \triangleq 2^{-\mathcal{H}_\mu(\mathbf{c})} \tag{11}$$

$$\mathbf{d}_\mu(\mathbf{c}_1,\mathbf{c}_2) \triangleq \Psi_\mu(\mathbf{c}_1 \sqcap \mathbf{c}_2). \tag{12}$$

Again, consider $\mu(e) \equiv \mathbf{1}_E$; denote as $\mathcal{H}(\bullet)$, $\Psi(\bullet)$ and $\mathbf{d}(\bullet,\bullet)$ the associated height, conciseness and pre-distance. We observe for this special case:

Lemma 1. For all $\mathbf{c} \in \mathcal{C}$,

$$\mathcal{H}(\mathbf{c}) = \infty \Rightarrow \mathbf{c} \in \Omega. \tag{13}$$

Proof: Assume $\mathbf{c} \notin \Omega$, and let $e \in E\backslash\mathbf{c}$ such that there is no $e' \in \mathbf{c}$ such that $e'\#e$, and let $n \triangleq \mathcal{H}([e'])$. Then $\mathcal{H}(\mathbf{c}) \leqslant n < \infty$ by definition of $\mathcal{H}(\bullet)$. □

As noted above, $\mathcal{H}_\mu(\bullet)$ and thus all functions derived from it, are invariant under isomorphisms. They thus lift without any further effort to functions on $\mathcal{C}_\lambda(\mathcal{E})$ instead of $\mathcal{C}(\mathcal{E})$; we will abuse of notation by using the same symbols for those lifted versions.

3. PETRI NETS AND THEIR SEMANTICS.

Petri Nets. We will turn now to an important subclass of event structures, obtained through Petri net models,

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

- (P, T, F) is a finite net, and
- $m : P \to \mathbb{N}$ 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 if $\forall p \in {}^\bullet t,\ m(p) > 0$. This enabled transition can *fire*, resulting in a new marking $m' = m - {}^\bullet t + t^\bullet$; this firing relation is denoted by $m[t\rangle m'$. A marking m is *reachable* if there exists a *firing sequence*, i.e. transitions $t_0 \dots t_n$ such that $m_0[t_0\rangle m_1[t_1\rangle \dots [t_n\rangle m$. A net is *safe* if for all reachable markings m, $m(p) \subseteq \{0,1\}$ for all $p \in P$.

Sequential semantics. The language $\mathcal{L}$ of $\mathcal{N}$ is the set of words $e_0 \dots e_n$ over a set E with a mapping $\lambda : E \to T$ such that $\lambda(e_0) \dots \lambda(e_n)$ is a firing sequence. Assume now that $\mathcal{L}$ is *trim*: any two words w, w' in $\mathcal{L}$ share their common prefix, i.e. if there are $u \in E^*, x, x' \in E^\infty$ and $e, e' \in E$ such that $w = uex$ and $w' = ue'x'$, then $\lambda(e) = \lambda(e')$ implies $e = e'$. The *sequential semantics* of $\mathcal{N}$ is given by event structure $\mathcal{E}_{seq} = (E, \leqslant_{seq}, \#_{seq}, \lambda)$, obtained from $\mathcal{L}$ by setting

(1) $e \leqslant_{seq} e'$ iff there exist $u, v \in E^*$ and $w \in E^\infty$ such that $uev e'w \in \mathcal{L}$, and
(2) $e \#_{seq} e'$ iff there exist $\bar{e}, \bar{e}' \in E$ and $u, v \in E^*$ such that $u\bar{e}, u\bar{e}' \in \mathcal{L}$ with $\lambda(\bar{e}) \neq \lambda(\bar{e}')$.

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

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

Definition 3. A net $ON = (B, E, G)$ 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 $x \in B \cup E$, the set $[x] = \{y \in B \cup E \mid y \leqslant_{ON} x\}$ 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.

The nodes of E are the *events*, those of B *conditions*. Occurrence nets are the mathematical form of the *partial order unfolding semantics* for Petri nets [(6)]; although more general applications are possible, we will focus here on unfoldings of *safe* Petri nets only.

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

- $h(P_1) \subseteq P_2$ and
- 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 $h(e_1) = h(e_2)$ then $e_1 = e_2$.

Branching processes $\beta_1 = (ON_1, \pi_1)$ and $\beta_2 = (ON_2, \pi_2)$ for $\mathcal{N}$ are isomorphic iff there exists a bijective homomorphism $h : ON_1 \to ON_2$ such that $\pi_1 = \pi_2 \circ h$. The unique (up to isomorphism) maximal branching process $\beta_\mathcal{U} = (ON_\mathcal{U}, \pi_\mathcal{U})$ of $\mathcal{N}$ is called the *unfolding* of $\mathcal{N}$; see [(6)] for a canonical algorithm to compute the unfolding of $\mathcal{N}$.

The **partial order semantics** for $\mathcal{N}$ is given by the event structure $\mathcal{E}_\mathcal{U} = (E_\mathcal{U}, \leqslant_\mathcal{U}, \#_\mathcal{U}, \pi^E_\mathcal{U})$ where $E_\mathcal{U}$ is the set of events in $\mathcal{N}$'s unfolding $\beta_\mathcal{U}$, and $\leqslant_\mathcal{U}$, $\#_\mathcal{U}$, and $\pi^E_\mathcal{U}$ are the restrictions to $E_\mathcal{U}$ of the corresponding elements of $\beta_\mathcal{U}$. By construction, the labeling $\pi^E_\mathcal{U}$ for $\mathcal{E}_\mathcal{U}$ is simple in the above sense: this property simply reflects the fact that no

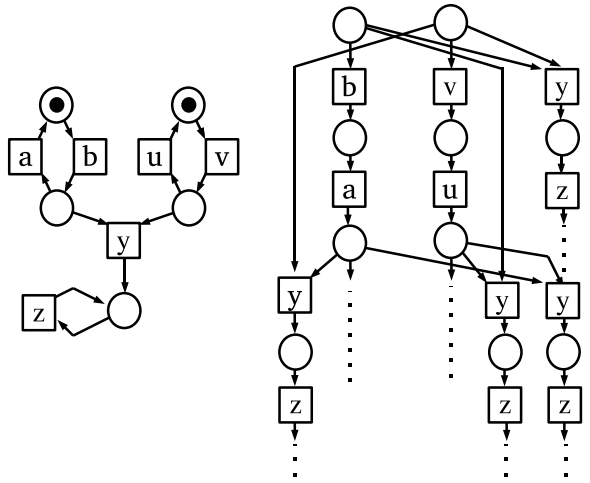

Fig. 1. Left: a Petri Net ; right: a prefix of its unfolding, with events bearing the name of π-image

transition can have more than one concurrent occurrence if the net is safe.

4. OBSERVABILITY AND DIAGNOSABILITY

Let $\eta : \mathbb{A} \to \mathbb{O}$ a partial mapping into an *observation alphabet* $\mathbb{O}$. For a given labeled prime event structure $\mathcal{E} = (E, \leqslant, \#, \lambda)$, let $E_\eta \triangleq \{e \mid \eta(\lambda(e)) \downarrow\}$ be the set of *visible* events, and $E_\varepsilon \triangleq \{e \mid \eta(\lambda(e)) \uparrow\}$ the set of *invisible* events. Using the above construction from the previous section, we obtain the *visible height* $\mathcal{H}_\eta(\bullet)$, *observable conciseness* $\Psi_\eta(\bullet)$ and pre-distance $\mathbf{d}_\eta(\bullet,\bullet)$, respectively, by setting $\mu \equiv \mathbf{1}_{E_\eta}$.

Observability. To avoid tedious case distinctions, we assume henceforth that all runs of $\mathcal{E}$ are of infinite height; if necessary, consider any finite-height run extended by an infinite chain of dummy events.

Definition 4. A labeled ES $(\mathcal{E}, \eta)$ is *observable* iff

$$\mathcal{H}(\mathbf{c}) = \infty \Rightarrow \mathcal{H}_\eta(\mathbf{c}) = \infty. \tag{14}$$

Topologies. Obviously, any choice of $\mu : \mathbb{A} \to \mathbb{R}_0^+$ and hence of $\mathbf{d}_\mu(\bullet,\bullet)$ defines a topology $\mathfrak{T}^\mu$, called the μ-topology, on Ω. Note that for $\mu \equiv \mathbf{1}_E$, we obtain the restriction - to Ω - of the Scott topology on $\mathcal{C}$; call this topology $\mathfrak{T}$. Further, denote as $\Omega_{/_\mu}$ the quotient space under $\mu \circ \lambda$-preserving isomorphism, with associated quotient topology $\mathfrak{T}_\mu$. In particular, set $\mathfrak{O} \triangleq \mathfrak{T}_\eta$.

Defining diagnosability. Let $\phi \in \mathbb{A}_\varepsilon$ be a *fault*. A configuration $\mathbf{c} \in \mathcal{C}(\mathcal{E})$ is called *faulty* iff $\mathbf{c} \cap \lambda^{-1}(\{\phi\}) \neq \emptyset$, and *healthy* otherwise. We can capture faultiness of configurations using $\mu \equiv \mathbf{1}_{\mathbb{A}_\eta}$, and then letting $\mathcal{H}_\phi(\mathbf{c}) \triangleq \mathcal{H}_{\mu_\phi}(\mathbf{c})$, etc. A configuration $\mathbf{c}$ is *faulty* iff $\mathcal{H}_\phi(\mathbf{c}) > 0$. Again, faultiness is invariant under isomorphism. Denote as Ω_F ($\mathcal{C}_F$) the set of faulty runs (configurations), and Ω_{NF} the set of healthy runs. We observe that if $\mathbf{c}$ is faulty, so is every extension of $\mathbf{c}$, i.e. every $\mathbf{c}' \in \mathcal{C}(\mathcal{E})$ such that $\mathbf{c} \subseteq \mathbf{c}'$ is faulty. As a consequence, we have:

Lemma 2. Ω_F is open in $\mathfrak{T}$.

Note, however, that Ω_F is in general neither open nor closed in $\mathfrak{O}$. We can distinguish three *diagnosis states*, given by sets of runs:

$$\begin{aligned} \textit{Fault} - \textit{definite} : FD &\triangleq \{\omega \in \Omega \mid [\omega]_\eta \subseteq \Omega_F\} \\ \textit{NF} - \textit{definite} : ND &\triangleq \{\omega \in \Omega \mid [\omega]_\eta \subseteq \Omega_{NF}\} \\ \textit{Indefinite} : ID &\triangleq \Omega \backslash (FD \cup ND)\,. \end{aligned}$$

It is of course not feasible to verify directly the *infinite* runs. In [(4)], a diagnoser system is built over *diagnoser states* that correspond to finite observation sequences : a diagnoser state represents the knowledge that can be derived about the eventual diagnosis, from a given finite observation. We shall not proceed here by constructing a diagnoser, since it is not feasible in general event structures; its state space would be infinite in general [2] . Rather, we give directly a definition of *eventual diagnosability* notions:

Definition 5. ϕ is *eventually F-diagnosable* for $(\mathcal{E}, \eta)$ iff Ω_F is open in $\mathfrak{O}$. Dually, ϕ is *eventually N-diagnosable* for $(\mathcal{E}, \eta)$ iff Ω_{NF} is open in $\mathfrak{O}$.

This is a notion that does not at all take the time after fault occurrence into account, contrary to e.g. [(27; 10)]. It generalizes the traditional definition from [(4)] given in the introduction, and the ones we presented for Petri nets in [(14; 16; 17)]. The corresponding *structural* characterization will be generalized in Theorem 3 below.

Metric characterization. Exploring the topology $\mathfrak{O}$ to characterize F-and NF-diagnosability shows us that both are equivalent, confirming corresponding results (see [(29)]) in the sequential case:

Theorem 1. If $(\mathcal{E}, \eta)$ is observable, then ϕ is eventually F-diagnosable for $(\mathcal{E}, \eta)$ iff for every faulty $\omega_\phi \in \Omega_F$, there exists a finite-height prefix $\mathbf{c}_\phi$ of ω_ϕ such that $\Omega_{\mathbf{c}_\phi} \subseteq \Omega_F$. Dually, if $(\mathcal{E}, \eta)$ is observable, then ϕ is eventually NF-diagnosable for $(\mathcal{E}, \eta)$ iff for every healthy $\omega_0 \in \Omega_{NF}$, there exists a finite prefix $\mathbf{c}_0$ of ω_0 such that $\Omega_{\mathbf{c}_0} \subseteq \Omega_{NF}$.

Proof: Fix ω_ϕ and assume ϕ is eventually F-diagnosable; then there exists $\delta = \delta(\omega_\phi)$ such that

$$\forall \omega \in \Omega_{NF} : \ \mathbf{d}_\eta(\omega_\phi, \omega) > \delta. \tag{15}$$

Let k be any integer such that $k > \log_2(\delta)$; then let $\mathbf{c}_\phi$ be the smallest prefix of ω_ϕ such that $\mathcal{H}_\eta(\mathbf{c}_\phi) = k$. By observability, $\mathcal{H}(\mathbf{c}) < +\infty$, and (15) implies that $\Omega_{\mathbf{c}_\phi} \subseteq \Omega_F$. The reverse implication is obvious. Finally, the proof for the characterization of NF-diagnosability is exactly analogous. □

We obtain the following additional result:

Theorem 2. If $(\mathcal{E}, \eta)$ is observable, then: ϕ is eventually NF-diagnosable for $(\mathcal{E}, \eta)$ iff it is eventually F-diagnosable for $(\mathcal{E}, \eta)$.

Proof: It suffices to exploit the symmetry of $\mathbf{d}_\eta(\bullet,\bullet)$ in the proof of Theorem 1. □

The astute reader will notice that a system may be diagnosable even without being observable as defined in

[2] Note that, for the case of Petri nets with sequential semantics (see below), the diagnoser construction is carried out in [(22)]

Def. 4. In the case of non-observability, all runs ω, ω' for which $\mathcal{H}_\lambda(\mathbf{c})$ is finite, satisfy $\mathbf{d}_\eta(\omega, \omega') = 0$. For ϕ to be F- or NF-diagnosable in $(\mathcal{E}, \eta)$, the runs of finite observable height must either all be faulty or all be healthy. In our view, this fact illustrates that all *interesting* diagnosis problems concern *observable* systems. - The equivalence of F-diagnosability and NF-diagnosability had been shown in [(29)] for the classical approach, using an enumeration argument that requires *sequential* semantics.

Structural Characterization. The following generalizes our results for unfoldings of safe Petri nets presented in [(14)]. For any two finite configurations $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}(\mathcal{E})$, say that $\mathbf{c}_2$ *dominates* $\mathbf{c}_1$, written $\mathbf{c}_1 \propto \mathbf{c}_2$, iff $\mathcal{E}_{\mathbf{c}_1}$ is a prefix of $\mathcal{E}_{\mathbf{c}_2}$. Event structure $\mathcal{E}$ is *quasi-regular* [3] iff for all $\mathbf{c} \in \mathcal{C}(\mathcal{E})$ of finite height, every (infinite) run $\omega \in \Omega(\mathcal{E}_{\mathbf{c}})$ contains at least one pair $(\mathbf{c}_1, \mathbf{c}_2)$ of finite height configurations as prefixes such that $\mathbf{c}_1 \subseteq \mathbf{c}_2$ and $\mathbf{c}_1 \propto \mathbf{c}_2$. In particular, all unfoldings of 1-safe Petri nets are quasi-regular: all infinite runs of these unfoldings must pass through an infinite number of finite configurations corresponding to the same net marking, since the number of reachable markings is finite. Any pair $(\mathbf{c}_1, \mathbf{c}_2)$ of such configurations with $\mathbf{c}_1 \subseteq \mathbf{c}_2$ satisfies $\mathbf{c}_1 \propto \mathbf{c}_2$ by construction of the unfolding. - To complete our preparations for Theorem 3, let $\mathbf{c} \sim_\eta \mathbf{c}'$ iff there is an η-isomorphism between $\mathbf{c}$ and $\mathbf{c}'$, and $\mathbf{c} \sim_\phi \mathbf{c}'$ iff $\mathbf{c}$ and $\mathbf{c}'$ are either both healthy or both faulty.

Theorem 3. If $(\mathcal{E}, \eta)$ is observable and quasi-regular, ϕ is eventually F-diagnosable for $(\mathcal{E}, \eta)$ iff for all configurations $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_1', \mathbf{c}_2' \in \mathcal{C}(\mathcal{E})$ of finite height such that

$$\mathbf{c}_1 \subseteq \mathbf{c}_1' \wedge \mathbf{c}_1 \propto \mathbf{c}_1'$$
$$\mathbf{c}_2 \subseteq \mathbf{c}_2' \wedge \mathbf{c}_2 \propto \mathbf{c}_2',$$

the following holds:

$$\left.\begin{array}{rc} & \mathbf{c}_1 \sim_\eta \mathbf{c}_2 \\ \wedge & \mathbf{c}_1' \sim_\eta \mathbf{c}_2' \\ \wedge & \mathcal{H}(\mathbf{c}_1) < \mathcal{H}(\mathbf{c}_1') \end{array}\right\} \Rightarrow \mathbf{c}_1' \sim_\phi \mathbf{c}_2'. \qquad (16)$$

Proof: To show the "if" part, assume $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_1', \mathbf{c}_2'$ violate (16), i.e. without loss of generality

(1) $\mathbf{c}_2'$ is faulty, but neither $\mathbf{c}_1'$ nor $\mathbf{c}_1$ are,
(2) for $i \in \{1, 2\}$, $\mathbf{c}_i' = \mathbf{c}_i \circ \mathbf{d}_i$, where $\mathbf{d}_i \in \mathcal{C}(\mathcal{E}_{\mathbf{c}_i})$ and $\mathbf{d}_1^1 \neq \emptyset$ ($\mathbf{d}_2$ may be empty) , and
(3) for $i \in \{1, 2\}$, $\mathbf{c}_i' \sim_\eta \mathbf{c}_i$ and $\mathbf{c}_i' \propto \mathbf{c}_i$.

It follows that there is a configuration $\mathbf{d}_i^2 \in \mathcal{C}(\mathcal{E}_{\mathbf{c}_i'})$ that is an isomorphic copy of $\mathbf{d}_i$. Iterating this argument, let $\mathbf{c}_i^1 \triangleq \mathbf{c}_i' = \mathbf{c}_1 \circ \mathbf{d}_i^1$ and $\mathbf{c}_i^{n+1} \triangleq \mathbf{c}_i^n \circ \mathbf{d}_i^{n+1}$ for $n \in \mathbb{N}$. Then by assumption, $\mathcal{H}(\mathbf{c}_1^n) \to_{n\to\infty} \infty$ (the same need not be true for the sequence of $\mathbf{c}_2^n$). We have $\mathbf{c}_i^n \sim_\eta \mathbf{c}_i$ for all n; by construction, all $\mathbf{c}_2^n$ are healthy, so ϕ can not be F-diagnosable for $(\mathcal{E}, \eta)$.

For "only if", suppose ϕ is not F-diagnosable for $(\mathcal{E}, \eta)$. Then there exists $\omega \in \Omega_F$ such that for any finite-height prefix $\mathbf{c}$ of ω, there is $\mathbf{c}' \in \mathcal{C}(\mathcal{E})$ that satisfies $\mathbf{c}' \sim_\eta \mathbf{c}$ and $\Omega_{\mathbf{c}'} \cap \Omega_{NF} \neq \emptyset$. But then one obtains a violation of (16) from the assumption that $\mathcal{E}$ is quasi-regular. □

[3] in reference to the more restrictive *regular* event structures that are the object of Thiagarajan's conjecture [(28)].

What Interleavings do and don't see. Figure 1 illustrates that choosing a partial order vs an interleaving semantics has important consequences. To see this, note that if the net behaviour is recorded in sequential form, we still have an event structure semantics; yet the resulting event structure is degenerate in the sense that **co** is empty. Defining metric topology etc. as above, let $\phi = \pi^{-1}(\{v\})$, and assume the observation labellings for $\mathcal{E}_{seq}$ and $\mathcal{E}_{\mathcal{U}}$ both satisfy $dom(\eta) = \pi^{-1}(\{a\})$. Then:

a) In sequential semantics, the net is not observable: the run $\omega_s \in \Omega(\mathcal{E}_{seq})$ which consists only of occurrences of u and v satisfies $\mathcal{H}_\eta(\omega_s) = 0$ and $\mathcal{H}_\lambda(\omega_s) = \infty$. Further, $(\mathcal{E}_{seq}, \eta)$ is neither F-diagnosable nor NF-diagnosable, since all runs without an occurrence y are observationally indiscernable from the run ω' formed only by occurrences of a and b; this $\sim_\eta$ class therefore contains both faulty and healthy runs.

b) However, with the same assumptions, $(\mathcal{E}_{\mathcal{U}}, \eta)$ is both observable and diagnosable; in fact, all runs $\omega \in \Omega(\mathcal{E}_{\mathcal{U}})$ are F-definite.

This example shows that, while our *framework* allows to choose a variety of different semantics, there are huge differences in whether or not a given Petri net is diagnosable, depending on the semantics.

5. CONCLUSION

Comparison with the classical approach. The definition for F-diagnosability given in Sampath, Lafortune et al [(27)] requires existence of a uniform bound on the "time" after occurrence of the fault before diagnosis.can be adapted to our framework - using a sequential event structure $\mathcal{E}$ obtained from a finite automaton - as follows: Let

$$\mathcal{C}_\phi^* \triangleq \{\mathbf{c} \in \mathcal{C}_F \mid \forall \mathbf{c}. \in \mathcal{C} : \ \mathbf{c}' \subseteq \mathbf{c} \Rightarrow \ \mathbf{c}' \notin \mathcal{C}_F\}$$

be the set of *minimal faulty configurations*. ϕ is *F-diagnosable* for $(\mathcal{E}, \eta)$ iff for every $\mathbf{c}_\phi \in \mathcal{C}_\phi^*$, there exists $K = K(\mathbf{c}) > 0$ such that the following holds: If $\mathbf{c} \in \mathcal{C}(\mathcal{E})$ is such that $\mathbf{c}_\phi$ is η-isomorphic to a prefix of $\mathbf{c}$, and the $\mathbf{1}$-height of $\mathbf{c}$ is bounded by K plus the height of $\mathbf{c}_\phi$, then $\mathbf{c}$ is also faulty:

$$\Psi_{\mathbf{1}}(\mathbf{c}_\phi) + K \leqslant \Psi_{\mathbf{1}}(\mathbf{c}) \Rightarrow \mathbf{c} \in \mathcal{C}_F. \qquad (17)$$

then $\mathbf{c}$ is also faulty. Note that this definition uses the $\mathbf{1}$-height, not observable height; under observability and for finite state systems, both are equivalent in the sense that the topologies obtained are the same; plugging $\Psi_\eta(\bullet)$ into (17) instead of $\Psi_{\mathbf{1}}(\bullet)$ will define the same systems as diagnosable. in fact, it suffices to adjust the value of K to the maximum of K and the number of states of the system. This definition had inspired the analogous one we have given in [(14)] for Petri nets, which are also finite state systems and therefore allow for the same uniform bound. In our setting here, which is more general and adequate for capturing *infinite* state systems as well, it is no longer feasible to use $\Psi_{\mathbf{1}}(\bullet)$ instead of $\Psi_\eta(\bullet)$; moreover, we believe it is preferable to base the definition of diagnosability on an accessible entity, namely the stream of *observations*, rather than unobservable system behaviour. The verification of diagnosability has been shown **PSPACE**-complete for

the sequential case in [(2)]. This carries over to the nonsequential case since eventual diagnosability for safe Petri nets is in **PSPACE**; in fact, it suffices in the worst case to compute $\mathcal{U}_\tau^{\mathbf{1}_E}$ for τ obtained as the number K of states of the net multiplied by the number of $\mathbb{A}$-isomorphism classes among the maximal configurations of $\mathcal{U}_K^{\mathbf{1}_E}$ (the proofs follows similar lines as a result in [(15)]).

Outlook: The topological framework presented here has the advantage of allowing for unified proofs, based on the properties of event structures regardless of the semantics that generates them. It is applicable to any kind of system model that has an event structure semantics, and potentially useful for capturing extensions such as incomplete models, or loss of alarm. Future work will address such extensions.

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.

REFERENCES

[1] P. Baldan, T. Chatain, S. Haar and B. König. Unfolding-based Diagnosis of Systems with an Evolving Topology. Proc. *CONCUR 2008*, *LNCS* **5201**, 203–217, Springer 2008.

[2] A. Bauer and S. Pinchinat. A topological Perspective on Diagnosis. *Proceedings 9th International Workshop on Discrete Event Systems WODES 2008.*

[3] P. Bonizzoni, G. Mauri, G. Pighizzini. About infinite traces, in: V. Diekert (Ed.), *Proc. ASMICS Workshop Free Partially Commutative Monoids*, Report TUM-I9002, TU München, 1990, pp 1–10.

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

[5] J. Engelfriet. *Branching Processes of Petri Nets.* Acta Informatica **28**:575–591, 1991.

[6] J. Esparza, S. Römer, and W. Vogler. An improvement of McMillan's unfolding algorithm. *Form. Meth. in Syst. Des.* **20**(3):285-310, 2002.

[7] E. Fabre and A. Benveniste. Partial Order Techniques for Distributed Discrete Event Systems: why you can't avoid using them. Discrete Event Dynamic Systems, 2007 (17), 355-403.

[8] E. Fabre, A. Benveniste, C. Jard, and S. Haar. Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach. *IEEE Trans. Aut. Control* **48**(5)Â·714–727, 2003.

[9] E. Fabre, A. Benveniste, C. Jard, and S. Haar. Distributed monitoring of concurrent and asynchronous systems. *Discrete Event Dynamic Systems* **15**(1):33–84, 2005

[10] S. Genc and S. Lafortune. Predictability of event occurrences in partially-observed discrete-event systems. *Automatica* **45**(2), pp 310–311, 2009.

[11] M.P. Cabasino, A. Giua, S. Lafortune, C. Seatzu. Diagnosability analysis of unbounded Petri nets. *Proc. of 48th IEEE Conference on Decision and Control (CDC),* 2009.

[12] M.P. Cabasino, A. Giua, C. Seatzu. Diagnosability of bounded Petri nets. *Proc. of 48th IEEE Conference on Decision and Control (CDC),* 2009.

[13] A. Giua and C. Xie. Control of safe ordinary Petri nets using unfolding. *Discrete Event Dynamic Systems* **15**(4):349–373, Dec. 2005.

[14] S. Haar, A. Benveniste, E. Fabre, and C. Jard. Partial Order Diagnosability of Discrete Event Systems Using Petri Net Unfoldings. *Proc. of 42nd IEEE Conference on Decision and Control (CDC),* 2003.

[15] S. Haar. Diagnosability and Branching Process Semantics. In: *Object Petri Nets, Processes, and Object Calculi*. Festschrift for R. Valk. Bericht (tech. report) **265**, pp.13–34, FB Informatik, Univ. Hamburg.

[16] S. Haar. Unfold and Cover: Qualitative Diagnosability for Petri Nets. Proc. 46th IEEE Conference on Decision and Control, 2007.

[17] S. Haar. Qualitative Diagnosability for Petri Nets. To appear in Proc. 48th IEEE Conference on Decision and Control, 2009.

[18] S. Haar. Types of Asynchronous Diagnosability and the *Reveals*-Relation in Occurrence Nets. To appear in: *IEEE Transactions on Automatic Control*; preliminary version at http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/haar-tac10.pdf.

[19] R. Kummetz and D. Kuske. The topology of Mazurkiewicz Traces. *Theor. Comp. Sci.* **305**:237–258, 2003.

[20] M.Z. Kwiatkowska. A Metric for Traces. *Inf. Proc. Lett.* **35**:129–135.

[21] F. Lin. Diagnosability of discrete event systems and its applications. *Discrete Event Dynamic Systems*, 4(1), 1994, pp. 197-212.

[22] A. Madalinski, F. Nouioua and P. Dague, "Diagnosability Verification with Petri Net Unfoldings". Proc. KES 2009. Extended version: *Rapport de recherche* (research report) 1516, UMR 8623, CNRS, Université Paris-Sud, , March 2009.

[23] K. McMillan. Using Unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. *4th CAV* pp. 164–174, 1992.

[24] T. Murata. Petri Nets: Properties, Analysis and Applications. *Proc. of the IEEE* vol. **77** no 4, April 1989.

[25] M. Nielsen, G. Plotkin, G. Winskel. Petri nets, event structures, and domains (I). *TCS* **13**:85–108, 1981.

[26] W. Reisig. *Petri nets.* Springer Verlag, 1985.

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

[28] P. S. Thiagarajan. Regular event structures and finite Petri nets: a conjecture. In: *Formal and Natural Computing*, LNCS 2300, pp. 244–253, Springer 2002.

[29] Y. Wang, S. Lafortune and Tae-Sic Yoo. Decentralized Diagnosis of Discrete Event Systems Using Unconditional and Conditional Decisions. Proc. 44th CDC, December 12-15, 2005

[30] G. Winskel. Event structures. *Advances in Petri nets,* LNCS **255**: 325–392, Springer Verlag, 1987.

[31] T. Yoo and S. Lafortune. Polynomial-Time Verification of Diagnosability of Partially-Observed Discrete-Event Systems. *IEEE Trans. Aut. Control* **47**(9):1491-1495 , 2002.