

What Topology tells us about Diagnosability in Partial Order Semantics

November 9, 2011

Stefan Haar
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

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 ? When this problem is investigated in the context of 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 *discrete 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 $\eta : T \rightarrow \mathbb{O}$ 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 $T^* \rightarrow \mathbb{O}^*$. A classical definition of diagnosability is given in [25]; we follow the equivalent presentation of [7]. Write $s \sim_\eta s'$ iff $s, s' \in T^*$ are mapped to the same observable word in \mathbb{O}^* , and call any sequence s such that ϕ occurs in s a *faulty* sequence, and all other sequences *healthy*. Then :

Definition 1 (Sequential Diagnosability) \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. Moreover, consider intrinsically asynchronous distributed systems, such as encountered in telecommunications or more generally in networked systems. Here, 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, we were led in [5] 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 [10]. 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. The approach was extended to graph transformation systems for modelling dynamically evolving system topologies in [3]. We have provided a series of results [17, 14, 15, 16] 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 *weak* observability and diagnosability properties in [17, 16].

Bauer and Pinchinat [4] have given a topological view on diagnosability in terms of sequential languages. The present work develops a framework that includes both sequential and partial order semantics, retrieving and generalizing as a special case the results of [4] and showing connections between weak and strong properties. The key construction is that of suitable metrics on event structures. For this, we generalize a standard construction to be found in [6, 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 for the sequential *and* non-sequential viewpoints. That is, both sequential languages as in [7, 4] and the partial order semantics given in [9, 24] and used in [11, 16], associate event structures to a system; and the metric topology given here coincides, on the sequential semantics, with the Cantor topology used

in [4]. With these tools, the properties of weak and strong diagnosability from [17, 16] become different instances of a general property, *eventual diagnosability*, for labeled event structures. The difference between the weak and strong properties lies thus in the choice of *semantics* that produces the event structure model of behaviour for the system that is investigated.

Structure of the paper: We begin in Section 2. with the basic definitions for (labeled) event structures. The following Section 3. investigates partial observation and diagnosability, and develops the main general results of this paper. Section 4 specializes to safe Petri nets, and studies properties characterizing weakly diagnosable nets. We then conclude in Section 5. The present article relates work first presented in the conference paper [13], along with new results concerning the diagnosability of 1-safe Petri nets (section 4) and extended discussions.

2 Event Structures

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 $f : A \rightarrow B$ 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 $\text{dom}(f) \triangleq \{a \in A \mid f(a) \downarrow\}$, and the *image* of f is

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

We shall be using throughout this paper *prime event structures (PES)* following Winskel et al [24, 28], with particular attention to labeling. Fix some alphabet $\mathbb{A} \neq \emptyset$.

Definition 2 A (labeled) prime event structure (over alphabet \mathbb{A}) is a tuple $\mathcal{E} = (E, \leq, \#, \lambda)$, where

1. $E = \text{supp}(\mathcal{E})$ is the support, or set of events of \mathcal{E} ,
2. $\leq \subseteq E \times E$ is a partial order satisfying the property of finite causes, i.e. setting $[e] \triangleq \{e' \in E \mid e' \leq e\}$, one has

$$\forall e \in E : |[e]| < \infty, \tag{1}$$

3. $\# \subseteq E \times E$ an irreflexive symmetric conflict relation satisfying the property of conflict heredity, i.e.

$$\forall e, e', e'' \in E : e \# e' \wedge e' \leq e'' \Rightarrow e \# e'', \tag{2}$$

4. $\lambda : E \rightarrow \mathbb{A}$ is a total mapping called the labelling. Events $e, e' \in E$ are concurrent, written $e \mathbf{co} e'$, iff neither $e \leq e'$ nor $e' < e$ nor $e \# e'$ hold.

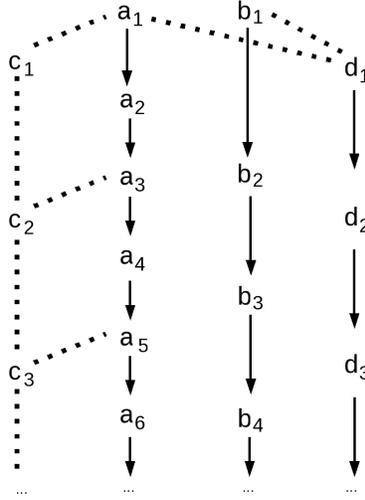


Figure 1: The simple event structure of Example 1. Arrows represent causal precedence \leq , and dashed lines stand for conflict $\#$; only minimal relations are represented, all others are generated by transitivity and inheritance.

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 simple¹ iff no label can occur concurrently on two different events; that is, iff

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

A simple labeled event structure will be called an SES.

Let $\mathcal{E}_1 = (E_1, \leq_1, \#_1, \lambda_1)$ and $\mathcal{E}_2 = (E_2, \leq_2, \#_2, \lambda_2)$ be two \mathbb{A} -labeled event structures. If (i) $E_1 \subseteq E_2$ and (ii) for all $e, e' \in E_1$,

$$e \#_1 e' \Leftrightarrow e \#_2 e' \quad \text{and} \quad e \leq_1 e' \Leftrightarrow e \leq_2 e',$$

then \mathcal{E}_1 is a sub-event structure of \mathcal{E}_2 .

Example 1. Consider the event structure $\mathcal{E} = (E, \leq, \#, \lambda_p)$ from Figure 1, where $E \triangleq \{a_i, b_i, c_i, d_i \mid i \in \mathbb{N}\}$ and $\mathbb{A} \triangleq \{a, a^*, b, b^*, c, c^*, d, d^*\}$, and let

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

$\lambda_p : E \rightarrow \mathbb{A}$ be defined by setting, for all $i \in \mathbb{N}$,

$$\begin{aligned}\lambda_p(a, 2i) &= a \quad \wedge \quad \lambda_p(a, 2i+1) = a^* \\ \lambda_p(b, 2i) &= b \quad \wedge \quad \lambda_p(b, 2i+1) = b^* \\ \lambda_p(c, 2i) &= c \quad \wedge \quad \lambda_p(c, 2i+1) = c^* \\ \lambda_p(d, 2i) &= d \quad \wedge \quad \lambda_p(d, 2i+1) = d^*.\end{aligned}$$

Define sets $A \triangleq \lambda_p^{-1}(\{a\})$, $A^* \triangleq \lambda_p^{-1}(\{a^*\})$, $\bar{A} \triangleq A \cup A^*$ and analogously $B, B^*, \bar{B}, C, C^*, \bar{C}, D, D^*, \bar{D}$. Let

1. for $i < j$, $a_i < a_j$, $b_i < b_j$ and $d_i < d_j$, but $c_i \# c_j$,
2. $a_{2i+1} \# c_i$, $a_i \# d_j$ and $b_i \# d_j$ for any $i, j \in \mathbb{N}$;

an illustration is given by Figure 1. One easily checks that $\mathcal{E} = (E, \leq, \#, \lambda)$ thus defined is an SES.

Prefixes and Configurations. The *set of causes* or *prime configuration* of $e \in E$ is $[e] \triangleq \{e' \mid e' \leq e\}$, as defined above. A *prefix* of \mathcal{E} is any downward closed subset $D \subseteq E$, i.e. such that for every $e \in D$, $[e] \subseteq D$. 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 $\mathcal{D}(\mathcal{E})$. Prefix $\mathbf{c} \in \mathcal{D}(\mathcal{E})$ 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. 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.

In the context of Example 1, one checks that, e.g., $[c_i] \cup [b_j]$ and $[a_i] \cup [b_j]$ are some of the configurations for all $i, j \in \mathbb{N}$; the runs are $\omega_{AB} \triangleq A \cup B$, $\omega_{c_i B} \triangleq [c_i] \cup B$ for $i \in \mathbb{N}$, and $\omega_D \triangleq D$.

2.1 Labeled event structure morphisms

Capturing partial observability follows the same idea as [7, 25]: a partial mapping associates a label to all observable events, and is undefined on the unobservable ones. Now, mappings for event structures require some additional conditions; we introduce here a dedicated class of morphisms for labeled event structures, which specializes Winskel's morphisms for event structures (see [28, 2]):

Definition 3 *Let $\mathcal{E}_1 = (E_1, \leq_1, \#_1, \lambda_1)$ and $\mathcal{E}_2 = (E_2, \leq_2, \#_2, \lambda_2)$ be two prime event structures. A partial mapping $f : E_1 \rightarrow E_2$ is a morphism iff for all $e_1 \in \text{dom}(f)$,*

1. $[f(e_1)] \subseteq f([e_1])$,
2. and for all $e'_1 \in \text{dom}(f)$,
 - (a) $f(e_1) \#_2 f(e'_1)$ implies $e_1 \#_1 e'_1$, and

(b) $f(e_1) = f(e'_1)$ and $e_1 \neq e'_1$ together imply that $e_1 \#_1 e'_1$.

A morphism $f : E_1 \rightarrow E_2$ is called an $(\mathbb{A}-)$ morphism iff, in addition,

1. $\text{dom}(\lambda_1) \subseteq \text{dom}(f)$ and $f(E_1) \subseteq \text{dom}(\lambda_2)$,
2. $\forall e \in E_1 : \lambda_1(e) = \lambda_2(f(e))$.

\mathcal{E}_1 and \mathcal{E}_2 are $(\mathbb{A}-)$ isomorphic, written $\mathcal{E}_1 \sim_{\mathbb{A}} \mathcal{E}_2$, iff there exist morphisms $f : E_1 \rightarrow E_2$ and $f^{-1} : E_2 \rightarrow E_1$ such that for all $e_1 \in \text{dom}(f)$ and all $e_2 \in \text{dom}(f^{-1})$,

$$f^{-1}(f(e_1)) = e_1 \quad \text{and} \quad f(f^{-1}(e_2)) = e_2.$$

That is, \mathbb{A} -morphisms

- may fuse conflicting events, but never any pair e_1, e_2 such that $e_1 < e_2$ or $e_1 \mathbf{co} e_2$, and
- have to respect labeling, i.e. two events mapped to the same target must either bear the same label or both be unobservable.

Note that Abbes [1] defines a different class of morphisms: *full* mapping $f : E_1 \rightarrow E_2$ is a morphism iff it is order-preserving between the underlying posets and if moreover f reflects conflict. This class is less appropriate than the above for our purposes since it does not allow for fusion of observationally equivalent conflicting configurations, nor for unobservable events.

For a given configuration $\mathbf{c} \in \mathcal{C}(\mathcal{E})$, we denote the set of configurations in \mathcal{E} that are \mathbb{A} -isomorphic images of \mathbf{c} as

$$[[\mathbf{c}]] \triangleq \{\mathbf{c}' \in \mathcal{C} \mid \mathbf{c}' \sim_{\mathbb{A}} \mathbf{c}\};$$

whenever the notation $[[\mathbf{c}]]_g$ is used below, it denotes analogously the isomorphism class w.r.t. a labeling g . The $[[\bullet]]$ -quotient of $\mathcal{C}(\mathcal{E})$,

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

consists of all possible *observation streams* under the observation mask λ . These quotient structures are central to the characterization of diagnosability given below.

For two prefixes D_1 and D_2 , write $D_1 \sqsubseteq_{\mathbb{A}} D_2$ iff D_1 is \mathbb{A} -isomorphic to a prefix of D_2 . For configurations $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}(\mathcal{E})$, let

$$[[\mathbf{c}_1]]_{\mathbb{A}} \sqcap [[\mathbf{c}_2]]_{\mathbb{A}} \triangleq [[\mathbf{c}_3]]_{\mathbb{A}},$$

where \mathbf{c}_3 is the unique \sqsubseteq -maximal prefix of \mathbf{c}_1 such that $\mathbf{c}_3 \sqsubseteq_{\mathbb{A}} \mathbf{c}_2$. This symmetric operation can be seen as the *intersection* of two configurations *up to \mathbb{A} -isomorphism*.

2.2 Pseudometrics.

We aim at giving the sets $\mathcal{C}(\mathcal{E})$ and $\Omega(\mathcal{E})$, and their observable images, *topological* structures that allow to reason about their properties. Canonical approaches to do so include Lawson or Scott topologies, *pseudometrics*; we will follow and generalize the latter approach, which is described e.g. in [19].

Recall that a *topology* on a set S is a collection $Top \subseteq 2^S$ of subsets of S , such that Top satisfies

1. $\emptyset \in Top$ and $S \in Top$,
2. if $A_1, \dots, A_n \in Top$, then $\bigcup_{i=1}^n A_i \in Top$, and
3. for arbitrary collections $(A_i)_{i \in I}$ of sets in Top , $\bigcup_{i \in I} A_i \in Top$;

the sets in Top are the *open sets* of the topological space (S, Top) , and $B \subseteq S$ is *closed* in (S, Top) iff $(S \setminus B) \in Top$.

An important class of topologies are those induced by *pseudometrics*: A mapping $d : S \times S \rightarrow \mathbb{R}$ such that for all $x, y, z \in S$,

1. $d(x, x) = 0$ and
2. $d(x, y) \leq d(x, z) + d(z, y)$ (*Triangle inequality*),

is called a *pseudometric* on S . Note that the triangle inequality implies that a pseudometric d satisfies, for all $x, y \in S$, both $d(x, y) \geq 0$ and $d(x, y) = d(y, x)$. A pseudometric d such that $d(x, y) = 0$ implies $x = y$ is a *metric*. From any pseudometric d on S , the d -induced topology is the smallest topology on S that contains all δ -balls w.r.t. d , i.e. all the sets

$$B_d(x, \delta) \triangleq \{y \in S \mid d(x, y) < \delta\}$$

for arbitrary $x \in S$ and $\delta \in (0, \infty)$.

The particularity of our pseudometrics defined below is that they are based on a weight function that allows to capture phenomena such as "invisibility" in partial observation. Our principal tool are μ -**Heights**: Let $\mu_{\mathbb{A}} : \mathbb{A} \rightarrow \mathbb{R}_0^+$ be any total mapping, and let $\mu \triangleq \mu_{\mathbb{A}} \circ \lambda$; we shall refer to μ as a *weight* function. As a particular case, consider $\mu(e) \equiv \mathbf{1}_E$: we will refer to 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 [6], see [19], when the counting weight is chosen; other choices of weights allow to generalize to observation and fault equivalence.

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

$$\mathcal{H}_{\mu}^*(\emptyset) \triangleq 0 \tag{4}$$

$$\mathcal{H}_{\mu}^*([e]) \triangleq \mathcal{H}_{\mu}^*([e] \setminus \{e\}) + \mu(e) \tag{5}$$

$$\mathcal{H}_{\mu}^*(D) \triangleq \sup_{e \in D} (\mathcal{H}_{\mu}^*([e])). \tag{6}$$

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

$$\mathcal{U}_\tau^\mu \triangleq \bigcup \{D \in \mathcal{D}(\mathcal{E}) \mid \mathcal{H}_\mu^*(D) \leq \tau\}, \quad (7)$$

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

$$\mathcal{H}_\mu(\mathbf{c}) \triangleq \sup\{\tau \mid \mathbf{c} \in \Omega(\mathcal{E}_\tau^\mu)\}. \quad (8)$$

This function gives the *lower height* of any configuration (w.r.t. μ); we call $\mathcal{H}_\mu^*(\mathbf{c})$ the *upper height* of \mathbf{c} . In fact, the terminology reflects that for any configuration \mathbf{c} , one has by construction

$$\mathcal{H}_\mu(\mathbf{c}) \leq \mathcal{H}_\mu^*(\mathbf{c}). \quad (9)$$

In figure 1, one has for the configuration $\mathbf{c}_1 \triangleq \{a_3, a_2, a_1, b_2, b_1\}$ that

$$\mathcal{H}_\mu(\mathbf{c}_1) = 2 < \mathcal{H}_\mu^*(\mathbf{c}_1) = 3,$$

while for $\mathbf{c}_2 \triangleq \mathbf{c}_1 \setminus \{a_3\}$,

$$\mathcal{H}_\mu(\mathbf{c}_2) = \mathcal{H}_\mu^*(\mathbf{c}_2) = 2. \quad (10)$$

We will call any configuration such that (10) holds *progressive*. In progressive configurations, all local processes advance at the same pace (unless blocked by conflicts), whereas general configurations such as \mathbf{c}_1 allow parts of the system to move ahead of the rest.

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

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

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

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. In this special case, $\mathbf{d}(\bullet, \bullet)$ is a *metric*; we further observe:

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

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

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

As noted above, $\mathcal{H}_\mu(\bullet)$ - and thus all the above functions derived from it - are invariant under isomorphisms.

Topology. The induced topology for the "counting" pseudometric $\mathbf{d}_{1_E}(\bullet, \bullet)$ has a base formed by the sets

$$\Omega_{\mathbf{c}} \triangleq \{\omega \in \Omega \mid \mathbf{c} \subseteq \omega\},$$

where \mathbf{c} is any *finite* configuration of \mathcal{E} , which are easily verified to be open. The set $\{\omega_{AB}\}$ - with the configuration AB given by $A \cup B$ - coincides e.g. with $\Omega_{\mathbf{c}_{31}}$, where $\mathbf{c}_{31} \triangleq [(a, 3)] \cup [(b, 1)]$. One obtains that $\{\omega_{AB}\}$, $\{\omega_D\}$ and all $\{\omega_{c_i B}\}$ are open; so are of course their unions and intersections. In particular, $\Omega_B = \{\omega_{AB}, \omega_{c_1 B}, \omega_{c_2 B}, \dots\}$ is also an open set. It should be noted that $\Omega_{\mathbf{c}}$ is in general *not* an open set, see below : in Figure 1, \bar{A} is a configuration that is infinite (though not maximal), and $\Omega_{\bar{A}}$ does not contain the open ball $B_{\mathbf{d}}(\bar{A}, 0.1)$. Taking another example, not for the configuration $A_2 = A \cup \{b_1, b_2\}$, $\Omega_{A_2} = \{\Omega_{AB}\}$ is *not* an open set, since any open neighbourhood of ω_{AB} must contain some $\omega_{c_i B} \in \Omega_{c_i B}$.

Further, one checks that configurations $[a_2] \cup [b_2]$ and $[c_2] \cup [b_4]$ are progressive, but e.g. $[a_6] \cup [b_4]$ is not.

3 Observability and Diagnosability

Let $\mathcal{E} = (E, \leq, \#, \lambda)$ with $\lambda : E \rightarrow \mathbb{A}$, and $\eta : \mathbb{A} \rightarrow \mathbb{O}$ a partial *observation mapping* into an *observation alphabet* \mathbb{O} . For a given labeled prime event structure, 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, 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 $\mu \equiv \mathbf{1}_{E_\eta}$. Write $\mathcal{E}_1 \sim_\eta \mathcal{E}_2$ iff the two structures with λ replaced by $\eta \circ \lambda$ are \mathbb{O} -isomorphic.

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} is eventually observable w.r.t. η iff

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

For an illustration, let $\mathbb{O} = \{a\}$ and define - in the context of Example 1 - the partial mapping $\eta : \mathbb{A} \rightarrow \mathbb{O}$ such that η maps a to a and is undefined otherwise. Then \mathcal{E} is not eventually observable w.r.t. η since one has, for every $i \in \mathbb{N}$,

$$\mathcal{H}(\omega_{c_i B}) = \infty \quad \wedge \quad \mathcal{H}_\eta(\omega_{c_i B}) = i - 1.$$

The traditional observability requires the existence of $k \in \mathbb{N}$ such that there are never more than k unobservable events in a row before the occurrence of an observable one. For safe Petri nets, this is equivalent to the system allowing no unobservable cyclic behaviour. Clearly, observable systems correspond to

eventually observable event structures; the converse is not true : in Figure 1, assume we have \odot the identity on \mathbb{A} , and $\lambda : E \rightarrow \mathbb{A}$ given by, for $i \in \mathbb{N}$,

$$\begin{aligned}\lambda(b_i) &= b \\ \lambda(c_i) &= c \\ \lambda(d_i) &= d \\ \lambda(a_i) &= \begin{cases} a & \text{if } \exists j \in \mathbb{N} : i = j^2 \\ \uparrow & \text{otherwise.} \end{cases}\end{aligned}$$

The resulting system is eventually observable, but not observable.

Topologies. Clearly, any choice of $\mu : \mathbb{A} \rightarrow \mathbb{R}_0^+$ and hence of $\mathbf{d}_\mu(\bullet, \bullet)$ defines a topology \mathfrak{T}^μ , 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

$$\begin{aligned}\mathcal{C}_{/\mu}(\mathcal{E}) &\triangleq \{[[\mathbf{c}]]_\eta \mid \mathbf{c} \in \mathcal{C}(\mathcal{E})\} \\ \Omega_{/\mu}(\mathcal{E}) &\triangleq \{[[\mathbf{c}]]_\eta \mid \mathbf{c} \in \mathcal{C}(\mathcal{E})\}\end{aligned}$$

the quotient spaces of configurations and runs, respectively, under $\mu \circ \lambda$ -preserving isomorphism, with associated quotient topology \mathfrak{T}_μ on $\Omega_{/\mu} = \Omega_{/\mu}(\mathcal{E})$. In particular, set $\mathfrak{D} \triangleq \mathfrak{T}_\eta$.

Defining diagnosability. Let $\Phi \subseteq E$ be a set of *invisible fault events*; in particular, no event in Φ is observable, i.e. $\lambda(\Phi) \cup \text{dom}(\eta) = \emptyset$. A configuration $\mathbf{c} \in \mathcal{C}(\mathcal{E})$ is called *faulty* iff $\mathbf{c} \cap \Phi \neq \emptyset$, and *healthy* otherwise. 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{D} .

We can distinguish three *diagnosis states*, given by sets of runs:

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

If the system is in state FD (or ND or ID), this means that its current configuration \mathbf{c} is such that

$$\Omega_{\mathbf{c}} \triangleq \{\omega \in \Omega \mid \mathbf{c} \subseteq \omega\} \subseteq FD(ND, ID)$$

It is of course not feasible to verify directly the *infinite* runs. In [7], 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 general². Rather, we give directly a definition of *eventual diagnosability*:

Definition 5 Φ is eventually F-diagnosable for (\mathcal{E}, η) iff Ω_F is open in \mathfrak{D} . Dually, Φ is eventually N-diagnosable for (\mathcal{E}, η) iff Ω_{NF} is open in \mathfrak{D} .

This is a notion that does not at all take the time after fault occurrence into account, contrary to e.g. [25, 12]. It generalizes the traditional definition from [7] given in the introduction, and the ones we presented for Petri nets in [17, 14, 15]. Traditional diagnosability notions such as the one [25, 12], or k-diagnosability (see e.g. [8] for a detailed discussion) impose a *uniform* bound on the logical time between occurrence and detection of a fault. Eventual diagnosability drops this constraint, thus becoming available in systems whose subsystems run in parallel with very weak synchronization or no synchronization at all.

An obvious objection to studying eventual diagnosability would be that any effective *positive* diagnosis of fault occurrence needs to know how long to wait for possible evidence. However, this real-time constraint is not coupled with the logical time constraints of Definition 1 or [8] in general. In fact, eventual diagnosability is *effective* diagnosability in systems where speedy progress is guaranteed, such that the effective latency until detection is uniformly bounded.

Metric characterization. Exploring the topology \mathfrak{D} to characterize F- and NF-diagnosability shows us that both are equivalent, confirming corresponding results (see [27]) in the sequential case:

Theorem 1 *If (\mathcal{E}, η) is eventually observable, then Φ is eventually F-diagnosable for (\mathcal{E}, η) iff for every faulty $\omega_\Phi \in \Omega_F$, there exists a finite-height prefix \mathbf{c}_Φ of ω_Φ such that $\Omega_{\mathbf{c}_\Phi} \subseteq \Omega_F$. Dually, if (\mathcal{E}, η) is eventually observable, then Φ is eventually NF-diagnosable for (\mathcal{E}, η) 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

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

Let k be any integer such that $k > \log_2(\delta)$; then let \mathbf{c}_Φ 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. \square

We obtain the following additional result:

²Note that, for the case of Petri nets with sequential semantics (see below), the diagnoser construction is carried out in [21]

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

Proof: Follows from the symmetry of $\mathbf{d}_\eta(\bullet, \bullet)$ in the proof of Theorem 1. \square

The astute reader will notice that a system may be eventually diagnosable even without being eventually observable as defined in 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 eventually F- or NF-diagnosable in (\mathcal{E}, η) , 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 - in the appropriate sense.

Note that equivalence of F-diagnosability and NF-diagnosability had been shown in [27] for the classical approach, using an enumeration argument that requires *sequential* semantics; the above generalization shows that it is an intrinsic, semantics-independent feature of diagnosis.

In the light of Theorem 2, we will henceforth drop the reference to F and NF and speak simply of *eventually diagnosable* labeled event structures.

Example. In the context of the event structure in Example 1, let us now choose $\mathbb{O} = \{b, d\}$ with $\text{dom}(\eta) = \{b, b^*, d, d^*\}$, where $\eta(b) = \eta(b^*) = b$ and $\eta(d) = \eta(d^*) = d$. If $\Phi \subseteq \{c_2, c_3, c_4, \dots\}$, then the net is not diagnosable since $\Omega_F = \bigcup_{i \in \mathbb{N}} \{\omega_{c_i B}, \omega_{c_i D}\}$ is not an open set in \mathfrak{D} ; any neighborhood of Ω_F in \mathfrak{D} contains $\omega_{AB} \in \Omega_{NF}$.

If one has, on the other hand, $\Phi \subseteq \overline{B}$, $\mathbb{O} = \{a, d\}$ and $\text{dom}(\eta) = \{a, a^*, d, d^*\}$, where $\eta(a) = \eta(a^*) = a$ and $\eta(d) = \eta(d^*) = d$, then \mathcal{E} is diagnosable with respect to η and Φ , since $\Omega_F = \{\omega_{c_i B} \mid i \in \mathbb{N}\} \cup \{\omega_{AB}\}$ is open in \mathfrak{D} .

Suffixes. 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. We have the following *suffix* objects: For $\mathbf{c} \in \mathcal{C}$ and $S \subseteq \mathcal{C}$, let

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

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

$$\begin{aligned} \mathcal{E}_{\mathbf{c}} &= (E_{\mathbf{c}}, \leq_{|E_{\mathbf{c}}}, \#_{|E_{\mathbf{c}}}, \lambda_{|E_{\mathbf{c}}}), \\ \text{where } E_{\mathbf{c}} &\triangleq \{e \in E \setminus \mathbf{c} \mid \forall e' \in \mathbf{c} : \neg(e \# e')\}, \end{aligned}$$

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 $\mathbf{c}'' \triangleq \mathbf{c} \cup \mathbf{c}' \in \mathcal{C}(\mathcal{E})$; 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} .

Structural Characterization. The characterization below result lifts the analogous one unfoldings of safe Petri nets presented in [17, 16] to *regular* event structures defined as follows: For any two finite configurations $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}(\mathcal{E})$, say that \mathbf{c}_2 *corresponds to* \mathbf{c}_1 , written $\mathbf{c}_1 \sim_{\mathcal{E}} \mathbf{c}_2$, iff $\mathcal{E}_{\mathbf{c}_1} \sim_{\Delta} \mathcal{E}_{\mathbf{c}_2}$. Clearly, $\sim_{\mathcal{E}}$ is an equivalence; event structure \mathcal{E} is *regular* iff it has a finite number of distinct $\sim_{\mathcal{E}}$ -classes. In particular, all unfoldings of 1-safe Petri nets are regular. In fact, all infinite runs of these unfoldings must pass through an infinite number of finite configurations corresponding to the behaviour after 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 $\mathbf{c}_1 \sim_{\mathcal{E}} \mathbf{c}_2$ by construction of the unfolding. The converse - can all regular event structures be constructed as unfoldings of 1-safe nets ? - is known as Thiagarajan's conjecture [26].

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}, η) is eventually observable and regular, Φ is eventually F-diagnosable for (\mathcal{E}, η) 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*

$$\begin{aligned} \mathbf{c}_1 \subseteq \mathbf{c}'_1 &\quad \wedge \quad \mathbf{c}_1 \sim_{\mathcal{E}} \mathbf{c}'_1 \\ \mathbf{c}_2 \subseteq \mathbf{c}'_2 &\quad \wedge \quad \mathbf{c}_2 \sim_{\mathcal{E}} \mathbf{c}'_2, \end{aligned}$$

the following holds:

$$\left. \begin{array}{l} \mathbf{c}_1 \sim_{\eta} \mathbf{c}_2 \\ \wedge \quad \mathbf{c}'_1 \sim_{\eta} \mathbf{c}'_2 \\ \wedge \quad \mathcal{H}(\mathbf{c}_1) < \mathcal{H}(\mathbf{c}'_1) \end{array} \right\} \Rightarrow \mathbf{c}'_1 \sim_{\Phi} \mathbf{c}'_2. \quad (16)$$

Proof: To show the “only 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 $\mathbf{c}'_i \sim_{\mathcal{E}} \mathbf{c}_i$.

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 $\mathbf{c}_i^1 \triangleq \mathbf{c}'_i = \mathbf{c}_i \circ \mathbf{d}_i^1$ 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, $\mathcal{H}(\mathbf{c}_1^n) \rightarrow_{n \rightarrow \infty} \infty$ (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}, η) .

For “if”, suppose Φ is not F-diagnosable for (\mathcal{E}, η) . 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 regular. \square

4 Application to Petri Nets

Petri Nets. We will turn now to an important instance of event structures, those linked to Petri net models.

Definition 6 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 \mathbb{N} . A **Petri net** is a tuple $\mathcal{N} = (P, T, F, m)$, where

- (P, T, F) is a finite net, and
- $m : P \rightarrow \mathbb{N}$ 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 2, 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 $m[t]m'$. A marking m is *reachable* if there exists a *firing sequence*, i.e. transitions $t_0 \dots t_n$ such that $m_0[t_0]m_1[t_1] \dots [t_n]m$. 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 $\lambda : E \rightarrow T$ 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 $\mathcal{E}_{seq} = (E, \leq_{seq}, \#_{seq}, \lambda)$, obtained from \mathcal{L} by setting

1. $e \leq_{seq} e'$ iff there exist $u, v \in E^*$ and $w \in E^\infty$ such that $ueve'w \in \mathcal{L}$, 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 $\lambda(\bar{e}) \neq \lambda(\bar{e}')$.

Partial order unfolding semantics. In a net $N = (P, T, F)$, let $<_N$ the transitive closure of F , and \leq_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

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

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

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

Definition 7 A net $ON = (B, E, G)$ is an **occurrence net** if and only if it satisfies

1. \leq_{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 $[x] = \{y \in B \cup E \mid y \leq_{ON} x\}$ 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 \leq_{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*. One notices quickly that complete occurrence nets form particular cases of event structures. The canonical association of an event structure to an occurrence net ON is by restricting \leq and $\#$ to the event set E , "forgetting" conditions. In particular, configurations of occurrence nets are defined as sets of events, i.e. configurations defined as above for the "stripped" event structure.

Occurrence nets are the mathematical form of the *partial order unfolding semantics* for Petri nets [18]; 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 $h : P_1 \cup T_1 \rightarrow P_2 \cup T_2$ 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 $h : ON_1 \rightarrow ON_2$ 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 [18] 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

$ON_{\mathcal{U}}$ spans a conflict free subnet $\mathbf{c}_{\mathcal{U}} = (E_{\mathbf{c}}, B_{\mathbf{c}}, G_{|(E_{\mathbf{c}} \times B_{\mathbf{c}}) \cup (B_{\mathbf{c}} \times E_{\mathbf{c}})})$ of $ON_{\mathcal{U}}$ by setting

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

The following results (see e.g. [18]) justify the use of unfoldings: The set $cut(\mathbf{c})$ of \leq -maximal nodes of $\mathbf{c}_{\mathcal{U}}$ 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 s ,

$$m(s) \triangleq \pi(s),$$

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, every reachable marking m of \mathcal{N} is reflected in this way by at least one configuration \mathbf{c} in $ON_{\mathcal{U}}$ such that $m(\mathbf{c}) = m$.

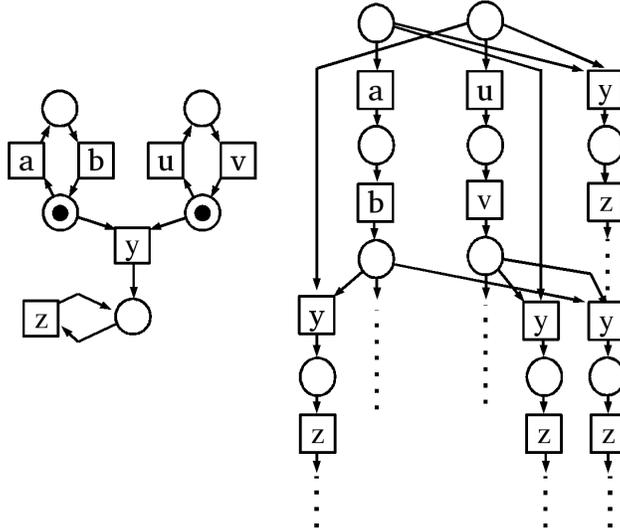


Figure 2: Left: a Petri Net ; right: a prefix of its unfolding, with events bearing the name of their π -image

The **partial order semantics** for \mathcal{N} is given by the event structure

$$\mathcal{E}_{\mathcal{U}} = (E_{\mathcal{U}}, \leq_{\mathcal{U}}, \#_{\mathcal{U}}, \pi_{\mathcal{U}}^E)$$

where $E_{\mathcal{U}}$ is the set of events in \mathcal{N} 's unfolding $\beta_{\mathcal{U}}$, and $\leq_{\mathcal{U}}$, $\#_{\mathcal{U}}$, and $\pi_{\mathcal{U}}^E$ are the restrictions to $E_{\mathcal{U}}$ of the corresponding elements of $\beta_{\mathcal{U}}$. By construction, the labeling $\pi_{\mathcal{U}}^E$ for $\mathcal{E}_{\mathcal{U}}$ is simple in the above sense: this property simply reflects the fact that no transition can have more than one concurrent occurrence if the net is safe.

Connecting the diagnosability notions. The notion of *F-diagnosability* given in Sampath, Lafortune et al [25] involves existence of a uniform bound on the “time” after occurrence of the fault before diagnosis. It can be adapted to our framework - using a sequential event structure \mathcal{E} obtained from a finite automaton - as follows: let

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

be the set of *minimal faulty configurations*. Φ is *F-diagnosable* for (\mathcal{E}, η) 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}_{Φ} is η -isomorphic to a prefix of \mathbf{c} , and the $\mathbf{1}$ -height of \mathbf{c} is bounded by K plus the $\mathbf{1}$ -height of \mathbf{c}_{Φ} , then \mathbf{c} is also faulty:

$$\mathcal{H}_1(\mathbf{c}_{\Phi}) + K \leq \mathcal{H}_1(\mathbf{c}) \Rightarrow \mathbf{c} \in \mathcal{C}_F. \quad (17)$$

then \mathbf{c} is also faulty. Note that this definition uses the $\mathbf{1}$ -height, not observable height; we will see below that, under observability, both are equivalent.

Characterizing diagnosable Petri nets. This definition had inspired the analogous one we have given in [17, 16] for safe Petri nets.

Definition 8 Let $\mathcal{N} = (P, T, F, m_0)$ a safe Petri net, $\eta : T \rightarrow \mathbb{O}$ a partial mapping, $\mathcal{U}_{\mathcal{N}} = (B, E, G, cut_0)$ its unfolding net, with labeling morphism $\lambda : E \rightarrow T$ given by the unfolding morphism. Let $\phi \in T \setminus \text{dom}(\eta)$ be a fault transition, and let $E_{\phi} \triangleq \lambda^{-1}(\phi)$. Denote by $\mathcal{C}_{\text{prog}}(\mathcal{N})$ the set of \mathcal{N} 's progressive configurations (compare (10)):

$$\mathcal{C}_{\text{prog}}(\mathcal{N}) \triangleq \{ \mathbf{c} \in \mathcal{C}(\mathcal{N}) \mid \mathcal{H}(\mathbf{c}) \leq \mathcal{H}_{\mu}^*(D) \}$$

We say that \mathcal{N} is **weakly observable w.r.t.** η iff its unfolding event structure $\mathcal{E}_{\mathcal{U}}$ is eventually observable w.r.t. η . A weakly observable (w.r.t. η) \mathcal{N} is **weakly diagnosable** w.r.t. η and ϕ iff there exists $n = n_{\mathcal{N}} \in \mathbb{N}$ such that for all configurations $\mathbf{c}_{\phi} \triangleq [e_{\phi}]$ with $e_{\phi} \in E_{\phi}$, every $\mathbf{c} \in \mathcal{C}_{\text{prog}}(\mathcal{N})$ such that

- (a) $\mathbf{c}_{\phi} \sqsubseteq \mathbf{c}$,
- (b) \mathbf{c} is not dead, and
- (c) $\mathcal{H}_1(\mathbf{c}) \geq \mathcal{H}_1(\mathbf{c}_{\phi}) + n$,

satisfies:

$$\forall \mathbf{c}' \in \mathcal{L} : \mathbf{c} \sqsubseteq_{\mathbb{O}} \mathbf{c}' \Rightarrow E_{\phi} \cap \mathbf{c}' \neq \emptyset. \quad (18)$$

Notice that the role of the set $\Phi \subseteq E$, which was arbitrary in the above study of diagnosability in *event structures*, is played here by the set E_ϕ of occurrences of the same transition ϕ . The definition implies that \mathcal{N} is weakly diagnosable w.r.t. ϕ and η iff $\mathcal{E}_{\mathcal{U}}(\mathcal{N})$ is diagnosable w.r.t. E_ϕ and η .

Let us first show the following auxiliary result:

Lemma 3 *If \mathcal{N} is observable, then there exists $n_\circ \in \mathbb{N}$ such that for any two configurations $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{C}(\mathcal{N})$ such that $\mathbf{c}_1 \sqsubseteq \mathbf{c}_2$ and $\mathbf{c}_1 \sim_\circ \mathbf{c}_2$, $\mathcal{H}(\mathbf{c}_2) \leq \mathcal{H}(\mathbf{c}_1)$.*

Proof: Suppose for every $n \in \mathbb{N}$ there exist $\mathbf{c}_1, \mathbf{c}_2$ such that $\mathcal{H}(\mathbf{c}_2) > \mathcal{H}(\mathbf{c}_1)$ while $\mathbf{c}_1 \sqsubseteq \mathbf{c}_2$ and $\mathbf{c}_1 \sim_\circ \mathbf{c}_2$. Then the pigeonhole principle implies, since the number of reachable markings of \mathcal{N} is bounded above by $2^{|P|}$, that for any $n > 2^{|P|}$, there exist $\mathbf{c}, \mathbf{c}' \in \mathcal{C}(\mathcal{N})$ such that

1. $m(\mathbf{c}) = m(\mathbf{c}')$
2. $\mathbf{c}_1 \sqsubseteq \mathbf{c} \sqsubseteq \mathbf{c}' \sqsubseteq \mathbf{c}_2$,
3. $\mathcal{H}(\mathbf{c}') \geq \mathcal{H}(\mathbf{c}) + 1$.

It follows that $\mathbf{c} \sim_\circ \mathbf{c}'$. Moreover, since $m(\mathbf{c}) = m(\mathbf{c}')$, any firing sequence leading from \mathbf{c} to \mathbf{c}' is again enabled in $m(\mathbf{c}')$, hence \mathcal{N} allows configurations $\mathbf{c}(n)$, $n \in \mathbb{N}$, such that $\mathbf{c} \sqsubseteq \mathbf{c}(1) \sqsubseteq \mathbf{c}(2) \sqsubseteq \dots$ and $\mathcal{H}(\mathbf{c}(n)) \geq \mathcal{H}(\mathbf{c}) + n$. This leads to a contradiction with weak observability as $n \rightarrow \infty$. \square

We then have:

Theorem 4 *Use the notations of Definition 8 and assume \mathcal{N} is weakly observable. Then \mathcal{N} is weakly diagnosable iff there exists $n \in \mathbb{N}$ such that for all $\mathbf{c}_\phi \in \mathcal{C}_\Phi(\mathcal{N})$ and $\mathbf{c} \in \mathcal{C}(\mathcal{N})$,*

$$\left. \begin{array}{l} \mathbf{c}_\phi \sqsubseteq \mathbf{c} \\ \mathbf{c} \text{ not dead} \\ \mathcal{H}_\circ(\mathbf{c}) \geq \mathcal{H}_\circ(\mathbf{c}_\phi) + n \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} \forall \omega \in \Omega(\mathcal{N}) : \\ (\mathbf{c} \sqsubseteq_\circ \omega) \Rightarrow \omega \in \Omega_F \end{array} \right. \quad (19)$$

Proof: Suppose first that \mathcal{N} is weakly diagnosable, i.e. $n_{\mathcal{N}}$ as in Definition 8 exists; then $n \triangleq \max(n_{\mathcal{N}}, n_\circ)$ with n_\circ from Lemma 3 has the above properties. Similarly, the existence of n as in the statement of the theorem implies that $n_{\mathcal{N}} \triangleq \max(n, n_\circ)$ satisfies the properties required in (19). \square

Example 2: What Interleavings do and don't see. Figure 2 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 \mathbf{co} is empty. Using the same definitions as above, let $\Phi = \pi^{-1}(\{v\})$, and assume the observation labellings for \mathcal{E}_{seq} and $\mathcal{E}_{\mathcal{U}}$ both satisfy $\text{dom}(\eta) = \pi^{-1}(\{a\})$. Then:

a) In sequential semantics, the net is *not* eventually 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_η class therefore contains both faulty and healthy runs.

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

The example allows to observe several important phenomena. In fact, it illustrates that decentralized systems with weak synchronization between sub-systems may elude diagnosis under the interleaved viewpoint, while being well captured under partial order semantics. In the example, consider now b the fault event, instead of v , and a observable. Then, the new system is neither 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.³ Thus, the partial order approach actually increases precision for partial observation of highly concurrent systems.

Another important observation here is that, as announced in the introduction, eventual diagnosability emerges as a fundamental property, for which weak and strong diagnosability are the concrete instances once the *semantics* has been chosen. In fact, it is often the choice of the semantics, i.e. of the mode of operation and observation for the system, that decides whether or not a given Petri net is eventually diagnosable.

5 Conclusion

We have cast the dynamics of discrete event systems in a general framework that allows to compare properties of the non-sequential and the sequential behaviour. On the level of abstraction granted by event structures, observability and diagnosability become general topological properties that specialize to existing concrete notions once the semantics (sequential or non-sequential) has been chosen. The verification of diagnosability has been shown to **PSPACE**-complete for the sequential case in [4]. This theoretical bound is a fortiori true for the non-sequential case. It is important now to develop efficient algorithms for verification of *weak* diagnosability; strong diagnosability has received treated in the existing literature, see e.g. [23, 22]). Current work is addressing these issues, based in particular on the result above (section 4) and an investigation of cutoff criteria for constructing suitable finite prefixes of unfoldings.

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; such extensions remain to be exploited.

³Thanks to A. Guia who made me discover this aspects by a remark in a workshop discussion

Future work will encompass *stochastic* models to make quantifiable predictions for weakly diagnosable systems.

Another approach to partial observation in concurrent systems, which has been introduced by the author in [14, 15, 16], 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.

An interesting extension of the work here consists in addressing infinite state systems with partial order methods as those presented here (as opposed to covering graphs as in [23]). More generally, developing methods for diagnosing event structures that are not regular is an interesting (and difficult) task.

Acknowledgments: This work was partly supported by the European Community's 7th Framework Programme under project DISC (*DI*stributed Supervisor Control of large plants), Grant Agreement INFSO-ICT-224498.

References

- [1] S. Abbes. A cartesian closed category of event structures with quotients. *Discrete Mathematics and Theoretical Computer Science*, 8(1):249–272, 2006.
- [2] P. Baldan, A. Corradini, and U. Montanari. Contextual petri nets, asymmetric event structures and processes. *Information and Computation*, 171(1):1–49, 2001.
- [3] 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.
- [4] Axel Bauer and Sophie Pinchinat. A topological perspective on diagnosis. In *9th International Workshop on Discrete Event Systems*, Gothenburg, Sweden, March 2008.
- [5] 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.
- [6] P. Bonizzoni, G. Mauri, and G. Pighizzini. About infinite traces. Report TUM-I9002, TU München, 1990.
- [7] C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, Boston etc, 1999.
- [8] Franck Cassez and Stavros Tripakis. Fault diagnosis with static and dynamic observers. *Fundam. Inform.*, 88(4):497–540, 2008.

- [9] J. Engelfriet. Branching Processes of Petri Nets. *Acta Informatica*, 28:575–591, 1991.
- [10] E. Fabre and A. Benveniste. Partial order techniques for distributed discrete event systems: why you can’t avoid using them. *Discrete Event Dynamic Systems: Theory and Applications*, 2007.
- [11] Éric Fabre, Albert Benveniste, Stefan Haar, and Claude Jard. Distributed monitoring of concurrent and asynchronous systems. *Discrete Event Dynamic Systems: Theory and Applications*, 15(1):33–84, March 2005.
- [12] S. Genc and S. Lafortune. Predictability of event occurrences in partially-observed discrete-event systems.
- [13] Stefan Haar. In *Proceedings of the 10th Workshop on Discrete Event Systems (WODES’10)*, Berlin, Germany, August.
- [14] Stefan Haar. Unfold and cover: Qualitative diagnosability for Petri nets. In *Proceedings of the 46th IEEE Conference on Decision and Control (CDC’07)*, pages 1886–1891, New Orleans, LA, USA, December 2007. IEEE Control System Society.
- [15] Stefan Haar. Qualitative diagnosability of labeled Petri nets revisited. In *Proceedings of the Joint 48th IEEE Conference on Decision and Control (CDC’09) and 28th Chinese Control Conference (CCC’09)*, pages 1248–1253, Shanghai, China, December 2009. IEEE Control System Society.
- [16] Stefan Haar. Types of asynchronous diagnosability and the *reveals*-relation in occurrence nets. *IEEE Transactions on Automatic Control*, 55(10):2310–2320, October 2010.
- [17] Stefan Haar, Albert Benveniste, Éric Fabre, and Claude Jard. Partial order diagnosability of discrete event systems using Petri net unfoldings. In *Proceedings of the 42nd IEEE Conference on Decision and Control (CDC’03)*, volume 4, pages 3748–3753, Hawaii, USA, December 2003. IEEE Control System Society.
- [18] S. Römer J. Esparza and W. Vogler. An improvement of mcmillan’s unfolding algorithm. *Formal Methods in System Design*, 20(3):285–310, 2002.
- [19] R. Kummetz and D. Kuske. The topology of Mazurkiewicz Traces. *Theoretical Computer Science*, 305:237–258, 2003.
- [20] M.Z. Kwiatkowska. A Metric for Traces. *Information Processing Letters*, 35:129–135, 1990.
- [21] Agnes Madalinski, Farid Nouioua, and Philippe Dague. Diagnosability verification with petri net unfoldings. *KES Journal*, 14(2):49–55, 2010.
- [22] C. Seatzu M.P. Cabasino, A. Giua. Diagnosability of bounded petri nets. In *Proc. of 48th IEEE Conference on Decision and Control (CDC)*, 2009.

- [23] S. Lafortune C. Seatzu M.P. Cabasino, A. Giua. Diagnosability analysis of unbounded petri nets. In *Proc. of 48th IEEE Conference on Decision and Control (CDC)*, 2009.
- [24] M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures, and domains (I). *Theoretical Computer Science*, 13:85–108, 1981.
- [25] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzi. Diagnosability of discrete-event systems. *IEEE Transactions on Automatic Control*, 40(9):1555–1575, 1995.
- [26] P. S. Thiagarajan. Regular event structures and finite petri nets: a conjecture. In *Formal and Natural Computing*, number 2300, pages 244–253. Springer, 2002.
- [27] Y. Wang, S. Lafortune, and Tae-Sic Yoo. Decentralized diagnosis of discrete event systems using unconditional and conditional decisions. In *Proc. 44th CDC*, 2005.
- [28] G. Winskel. Event structures. In *Advances in Petri nets*, number 255 in LNCS, pages 325–392. Springer Verlag.