Joint 48th IEEE Conference on Decision and Control and
28th Chinese Control Conference
Shanghai, P.R. China, December 16-18, 2009

Qualitative Diagnosability of labeled Petri Nets revisited

Stefan Haar
INRIA LSV
Ecole Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 CACHAN Cedex, France35042 Rennes cedex, France.
stefan.haar@inria.fr, haar@lsv.ens-cachan.fr

***Abstract*— In recent years, classical discrete event fault diagnosis techniques have been extended to Petri Net system models under partial order semantics [8], [9], [13]. In [14], we showed how to take further advantage of the partial order representation of concurrent processes, by decomposing the unfolding into 'facets', formed by subnets whose events either all occur eventually, or none of them occurs. A notion of *q(ualitative)*-diagnosability was proposed in [14] based on this decomposition. The present paper corrects the definition of q-diagnosability and develops its properties. Sufficient and necessary criteria, on the transition labeling, for q-diagnosability are shown; for their verification, and diagnosis itself, compact data structures are sufficient.**

I. Introduction

The diagnosis of large distributed systems requires to make good use of concurrency, rather than falling into its traps such as state space explosion due to enumeration of interleaved behaviours. Petri nets (see e.g. [16], [22], [23], [24], [17], [12], [11]) and their partial order unfoldings [20], [6], [18] have been increasingly used in recent years for both fault diagnosis ([8], [9], [13]) and control (see e.g. [15]) of asynchronous discrete event systems. The advantage of this semantics lies in the space reduction for representing nonsequential processes that have a high degree of parallelism. In unfoldings, sets of concurrent events are not ordered, which means they have to be represented only once (by one partial order) rather than by giving all their interleavings whose number is exponential in the size of the concurrent set. The gain in space therefore depends heavily on the degree of parallelism; the motivation is thus very strong in highly distributed systems such as telecommunication networks, see [9] and the discussion in the reference [7], entirely dedicated to the necessity of *true concurrency* in the study of distributed discrete event systems. In [8], [9], [13], fault diagnosis for a Petri net model $\mathcal{N}$ is performed by unfolding the labelled product of $\mathcal{N}$ and an observed alarm pattern $\mathcal{A}$, also in Petri net form. In [13], we have presented a characterization of diagnosability adapted to partial order semantics of 1-safe Petri nets. Subsequently, we studied in [14] an aspect of the relational structure of occurrence nets that is relevant for observation and diagnosis. Recall that occurrence nets carry a relational structure (known as *event structures* in the literature, see [21]) that consists of a partial order $<$ and a conflict relation $\#$; pairs that are neither ordered, nor in conflict, are collected in the complement relation $\mathbf{co}$ for *concurrency*. Within this structure, relation $a \triangleright b$ holds for events iff a *logically covers* b, or *leads to* b in the sense that whenever a occurs, b must eventually occur as well. *Facets* are subnets of the unfolding in which *any* two events cover one another. As a consequence, if some event in a facet occurs, eventually all other events of the facet have to occur in any fair execution (i.e. assuming progress: no enabled event remains enabled forever without occurring). In [14], we introduced the concept of $\mathbf{q}$-*diagnosability*, for *qualitative* diagnosability as opposed to quantitative criteria like those in [13], [25]; it is a property that is specific to partial order semantics, with no equivalent in the sequential case. The present paper focuses on $\mathbf{q}$-diagnosability, its properties and practical verification, as well as its relation with facets.

Overview: The paper is organized as follows: We start in Section II by recalling basic definitions concerning the basic tools and goals. In Section III, we recall the relation $\triangleright$ and the definition and properties of facets from [14]. Section IV is the heart of the article, studying properties of $\mathbf{q}$-diagnosability; Section V concludes.

II. Petri Nets, Unfoldings, and Diagnosis

Nets and homomorphisms: A *net* is a triple $N = (\mathcal{P}, \mathcal{T}, F)$, where $\mathcal{P}$ and $\mathcal{T}$ are disjoint sets of *places* and *transitions*, respectively, and $F \subseteq (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})$ is the *flow relation.* In figures, places are represented by circles, rectangular boxes represent transitions, and arrows represent F. Note that we consider only *ordinary* Petri nets here (compare [23]); that is, the weight of all arcs is equal to 1, and will be omitted. Let $<$ be the transitive closure of F and $\leqslant$ the reflexive closure of $<$. For node $x \in \mathcal{P} \cup \mathcal{T}$, call ${}^\bullet x \triangleq \{x' \mid F(x', x)\}$ the *preset* and $x^\bullet \triangleq \{x' \mid F(x, x')\}$ the *postset* of x; further, call $\lceil x \rceil \triangleq \{x' \mid x' < x\}$ the *prime configuration* or *cone* of x. A *net homomorphism* from N to N' is a map $\pi : \mathcal{P} \cup \mathcal{T} \longmapsto \mathcal{P}' \cup \mathcal{T}'$ such that:

1) $\pi(\mathcal{P}) \subseteq \mathcal{P}'$, $\pi(\mathcal{T}) \subseteq \mathcal{T}'$, and
2) $\pi_{|{}^\bullet t} : {}^\bullet t \to {}^\bullet \pi(t)$ and $\pi_{|t^\bullet} : t^\bullet \to \pi(t)^\bullet$ induce bijections, for every $t \in \mathcal{T}$.

Definition 1: Two nodes x, x' of a net N are *in conflict,* written $x \# x'$, if there exist transitions $t, t' \in \mathcal{T}$ such that

1) $t \neq t'$,
2) ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$, and

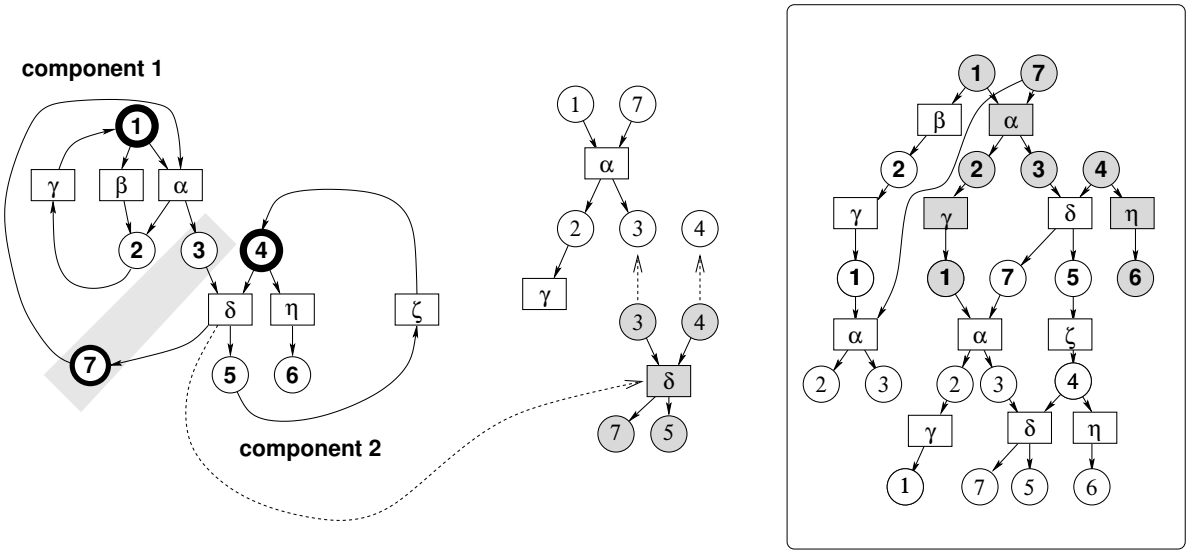

Fig. 1. Unfolding procedure in the context of the Petri net example from [8]

3) $t \leqslant x$ and $t' \leqslant x'$.

A node x is said to be in *self-conflict* iff $x\#x$. An *occurrence net* (ON) is a net $ON = (\mathcal{B}, \mathcal{E}, F, \mathbf{c}_0)$, with the elements of $\mathcal{B}$ called *conditions* and those of $\mathcal{E}$ *events*, such that:

1) no self-conflict: $\forall x \in \mathcal{B} \cup \mathcal{E} : \neg[x\#x]$;
2) $\leqslant$ is a partial order: $\forall x \in \mathcal{B} \cup \mathcal{E} : \neg[x < x]$;
3) $\forall x \in \mathcal{B} \cup \mathcal{E} : |\lceil x \rceil| < \infty$;
4) no backward branching: $\forall b \in \mathcal{B} : |{}^\bullet b| \leq 1$.
5) the set $\mathbf{c}_0 \triangleq \min(ON)$ of minimal nodes of ON is contained in $\mathcal{B}$.

Here, we add w.l.o.g. restriction 5.) for convenience; it is not required, e.g., in [5]. Nodes x and x' are *concurrent,* written x **co** x', if neither $x \leqslant x'$, nor $x' \leqslant x$, nor $x\#x'$ hold. A *co-set* is a set $\mathcal{X}$ of pairwise concurrent conditions. A maximal co-set $\mathcal{X}$ w.r.t. set inclusion is called a *cut*, and generically denoted by $\mathbf{c}$; in particular, $\mathbf{c}_0$ is a cut, called the *initial* cut of $\mathcal{ON}$. We note for future reference that occurrence nets are a special case of *event structures* [21]:

Definition 2: $(E, <, \#)$ is an *event structure (ES)* iff:

1) $(E, <)$ is a countable, partially ordered set,
2) $\lceil e \rceil$ is finite for all $e \in E$,
3) $\# \subseteq E \times E$ is symmetric and irreflexive, and such that $\forall\, x, y, z \in E$: $x\#y$ and $y < z$ together imply $x\#z$.

Petri Nets: Let $N = (\mathcal{P}, \mathcal{T}, F)$ be a finite net. A *marking* of net N is a multi-set $M \in \mathfrak{M}(\mathcal{P})$. A *Petri net* (PN) is a pair $\mathcal{N} = (N, M_0)$, with $M_0 \in \mathfrak{M}(\mathcal{P})$ an *initial* marking. $t \in \mathcal{T}$ is *enabled* at M, written $M \xrightarrow{t}$, if for all $p \in {}^\bullet t$, $M(p) \geqslant 1$. If $M \xrightarrow{t}$, then t can *fire*, leading to

$$M' \quad = \quad (M - 1_{{}^\bullet t}) + 1_{t^\bullet},$$

where symbol 1 denotes the set indicator function; write in that case $M \xrightarrow{t} M'$. The set $\mathbf{R}(M_0)$ contains the markings of $\mathcal{N}$ *reachable* through $\longrightarrow$. In the figures here, marked places are highlighted in thick; A Petri net $\mathcal{N} = (N, M_0)$ is *k-safe* if for all $M \in \mathbf{R}(M_0)$ and places p, $M(p) \leqslant k$. Only safe, i.e. 1-safe nets or *1-PNs* are considered here; their reachable markings are sets $M \subseteq \mathcal{P}$.

Branching Processes and Unfoldings: The branching process semantics reflects the partial order behavior of Petri nets in occurrence nets, thus allowing for structural analysis.

Definition 3: A branching process *of the safe Petri net* $\mathcal{N} = (\mathcal{P}, \mathcal{T}, F, M_0)$ *is given by a pair* $\pi = (\mathcal{ON}, \pi)$*, where* $\mathcal{ON} = (\mathcal{B}, \mathcal{E}, G, \mathbf{c}_0)$*, and* π *is a homomorphism from* $\mathcal{ON}$ *to* N*, such that:*

- $\pi(\mathbf{c}_0) = M_0$*;*
- *for all* $e, e' \in \mathcal{E}$, ${}^\bullet e = {}^\bullet e'$ *and* $\pi(e) = \pi(e')$ *together imply* $e = e'$.

Definition 4: Occurrence net $\rho = (\mathcal{B}_\rho, \mathcal{E}_\rho, G_\rho, \mathbf{c}_0(\rho))$ *is a* (structural) prefix *of* $\mathcal{ON}$*, written* $\rho \sqsubseteq \mathcal{ON}$*, iff*

1) $\mathcal{B}_\rho \subseteq \mathcal{B}$, $\mathcal{E}_\rho \subseteq \mathcal{E}$, *and* $G_\rho = G_{|(\mathcal{B}_\rho \times \mathcal{E}_\rho) \cup (\mathcal{E}_\rho \times \mathcal{B}_\rho)}$*;*
2) $e \in \mathcal{E}_\rho \Rightarrow {}^\bullet e \cup e^\bullet \subseteq \mathcal{B}_\rho$
3) $\mathbf{c}_0 = \mathbf{c}_0(\rho)$*; and*
4) ρ *is* causally closed*: if* $x' \leqslant x$ *and* $x \in \rho$*, then* $x' \in \rho$.

A prefix κ *of* $\mathcal{ON}$ *is a* configuration *if* κ *is* conflict-free*, i.e. no two nodes from* κ *are in conflict. A* maximal *configuration (w.r.t. set inclusion) is called a* run *and generically denoted* ω*; denote the set of runs as* Ω*. For* π, π' *two branching processes,* π' *is a* prefix *of* π*, written* $\pi' \sqsubseteq \pi$*, if there exists an injective homomorphism* ψ *from* $\mathcal{ON}'$ *into a prefix of* $\mathcal{ON}$*, such that* ψ *induces a bijection between the initial cuts* $\mathbf{c}_0$ *and* $\mathbf{c}'_0$*, and the composition* $\pi \circ \psi$ *coincides with* π'.

By theorem 23 of [5], there exists a unique (up to an isomorphism) $\sqsubseteq$-maximal branching process, called the *unfolding* of $\mathcal{N}$ and denoted $\mathcal{U}(\mathcal{N})$; by abuse of notation, we will also use $\mathcal{U}(\mathcal{N})$ for the occurrence net obtained by the unfolding. The principle for effectively constructing the unfolding (see [5], [8], [15]) is as follows: a copy of initial marking M_0 yields the initial conditions; events are appended to concurrent conditions that enable them, and are followed by the post-conditions they create. An illustration is given in Figure 1, taking up the running example from [8], [13]. Petri net $\mathcal{N}$ is

shown on the left, and a branching process $\pi = (ON, \pi)$ of $\mathcal{N}$ on the right hand side. Conditions are labeled by places, events by transitions. A configuration is shown in grey. The mechanism for constructing the unfolding of $\mathcal{N}$ is illustrated in the middle.

Configurations: Every *finite* configuration κ terminates at a cut, denoted $\mathbf{c}_\kappa$, such that $\kappa \mapsto \mathbf{c}_\kappa$ is bijective; for each cut $\mathbf{c}$, the *downward closure* $\downarrow(\mathbf{c}) \triangleq \{x \mid \exists b \in \mathbf{c} : x \leq b\}$ is the unique configuration κ such that $\mathbf{c} = \kappa_\mathbf{c}$. Conversely, every finite configuration κ corresponds to a unique reachable marking $M(\kappa)$ given by $M(\kappa) \triangleq \pi(\mathbf{c}_\kappa)$. We call configurations (and their cuts) that lead to the same marking *marking equivalent*, and write

$$\kappa \sim_M \kappa' \quad \textit{iff} \quad M(\kappa) = M(\kappa') \tag{1}$$

$$\mathbf{c} \sim_M \mathbf{c}' \quad \textit{iff} \quad \downarrow(\mathbf{c}) \sim_M \downarrow(\mathbf{c}'). \tag{2}$$

The nonsequential executions of safe Petri net $\mathcal{N}$ are in one-to-one correspondence with the configurations of $\mathcal{U}(\mathcal{N})$.

Finite Complete Prefix: If $\mathcal{U}(\mathcal{N})$ is infinite, we are naturally interested in finite prefixes of $\mathcal{U}(\mathcal{N})$ that are *complete* in the sense that their analysis allows to derive results for all of $\mathcal{U}(\mathcal{N})$. The definition and size of such prefixes varies with the intended purpose; see [18] for a systematic treatment. We use here the following Definition, similar to that in [15]:

Definition 5: The order 1 unfolding, denoted $\mathcal{U}_1(\mathcal{N})$, is a finite prefix of the unfolding obtained by stopping the construction of the unfolding when we reach a *cut-off* event e, i.e., an event such that:

- EITHER firing of $\lceil e \rceil$ brings back to the initial marking: $M(\lceil e \rceil) = M_0$;
- OR there exists another event e' with the following properties:
  1) The prime configuration for e' is a prefix of that of e: $\lceil e' \rceil \subseteq \lceil e \rceil$;
  2) the markings reached firing the two congurations are equivalent: $M(\lceil e \rceil) = M(\lceil e' \rceil)$.

In the following we call e' the *mirror transition* of e in $\tilde{\mathcal{N}}_1(M_0)$. Once we have constructed $\mathcal{U}_1(\mathcal{N})$, assume we continue the unfolding until we reach an event e such that there exist another event e' with the following properties:

- either e' does not belong to $\mathcal{U}_1(\mathcal{N})$ or it is a cut-off event of $\mathcal{U}_1(\mathcal{N})$;
- The prime configuration for e' is a prefix of that of e: $\lceil e' \rceil \subseteq \lceil e \rceil$;
- the two configurations are marking-equivalent: $M(\lceil e \rceil) = M(\lceil e' \rceil)$.

The resulting net, denoted $\mathcal{U}_2(\mathcal{N})$, is called *order 2 unfolding*; recursively, one obtains *order n unfoldings* $\mathcal{U}_n(\mathcal{N})$ for $n \geq 2$ in the same way.

Note that the initial definition from [20] used as cutoff criterion the *cardinality*, i.e. $|\lceil e' \rceil| < |\lceil e \rceil|$, which would lead to a shorter prefix in general yet not guarantee completeness w.r.t. computing $\triangleright$.

Diagnosis and Diagnosability: We assume the system to be diagnosed is modeled as a safe Petri net $\mathcal{N} = (N, M_0)$ with $N = (\mathcal{P}, \mathcal{T}, F)$, a set A of alarm labels, and an A-labeled alarm pattern $\mathcal{A}$ in the form of a conflict-free occurrence net; that is, the events of $\mathcal{A}$ are labeled by the observed alarms in A, and in the order they were observed. The conditions in the occurrence net $\mathcal{A}$ are dummy conditions whose only purpose is to exhibit precedence ordering.

When unfolding labeled nets, one obtains a labeled occurrence net; we denote with the same symbol λ the mapping $\mathcal{E} \to A$ obtained by $\lambda \circ \pi$, that is, an event e is labeled by the label $\lambda(t)$ of the transition $t \triangleq \pi(e)$. Denote as $\mathcal{I} \triangleq \lambda^{-1}(\{\varepsilon\})$ the set of *invisible* transitions, where $\lambda : \mathcal{T} \to A$ is the labeling function and $\varepsilon \in A$ the empty symbol. Dually, let $\mathcal{T}_A \triangleq \mathcal{T} \backslash \mathcal{I}$ be the set of *visible* transitions, and set $\mathcal{E}_\mathcal{I} \triangleq \pi^{-1}(\mathcal{I})$ and $\mathcal{E}_A \triangleq \pi^{-1}(\mathcal{T}_A)$. Let $\mathcal{U}(\mathcal{N}) = (\mathcal{B}, \mathcal{E}, G, \mathbf{c}_0)$ be the unfolding of $\mathcal{N}$ with homomorphism π, and denote as $O \triangleq \pi^{-1}(\mathcal{T}_A)$ the set of observable events. Further, let $\phi \in \mathcal{I}$ be a fault to be observed; let $\mathcal{E}_\phi \triangleq \pi^{-1}(\{\phi\})$. For configurations κ, κ' of $\mathcal{N}$, write $\kappa \sim_A \kappa'$ iff the sets κ_A, κ'_A of observable events of κ and κ', respectively, are isomorphic partially ordered sets (with the order relation induced by $\leqslant$). Let $\phi \in \mathcal{I}$ be a *fault*[1] to be diagnosed. Configurations κ, κ' are *ϕ-equivalent* iff either both contain a ϕ-event, or neither of them does:

$$\kappa \sim_\phi \kappa' \quad \textit{iff} \quad [\kappa \cap \mathcal{E}_\phi \neq \emptyset \iff \kappa' \cap \mathcal{E}_\phi \neq \emptyset.] \tag{3}$$

The asynchronous diagnosis of [8], [9], [13] proceeds as follows: Take the Petri net model $\mathcal{N}$ of the system, with transition labeling $\lambda : \mathcal{T} \to A$ taking values in an alphabet A of alarms, and the Petri net representation $\mathcal{A}$ of the observed alarm pattern. Then form the product net $\mathcal{N} \times \mathcal{A}$ by fusing transitions carrying the same label. All executions of $\mathcal{N} \times \mathcal{A}$ correspond to executions of $\mathcal{N}$; the converse is obviously not true. Moreover, not all executions of $\mathcal{N} \times \mathcal{A}$ cover all of $\mathcal{A}$; in general, only a proper prefix of the observation is explained by a given run of $\mathcal{N} \times \mathcal{A}$. In the unfolding $\mathcal{U}(\mathcal{N} \times \mathcal{A})$, take all those branches that *fully explain* $\mathcal{A}$; the corresponding executions of $\mathcal{N}$ form the *diagnosis set* of all possible explanations of $\mathcal{A}$ in the model $\mathcal{N}$.

Convergence of the diagnosis procedure involves the computation of the full unfolding $\mathcal{U}(\mathcal{N} \times \mathcal{A})$. That is, to be effective, it requires that $\mathcal{U}(\mathcal{N} \times \mathcal{A})$ be finite; this means we must require that $\mathcal{N}$ contains no *invisible cycles*. In other words, sufficiently many transitions of $\mathcal{N}$ must carry a non-empty label and thus be *visible*, such that the net cannot leave a marking M and then return to M without having produced a visible alarm on the record. That is (compare [13]), for any two configurations κ, κ' such that κ is a proper prefix of κ' and $\kappa \sim_M \kappa'$, there must be at least one visible event in $\kappa \backslash \kappa'$ for observability of the net. This property will be assumed throughout. *Diagnosability* is the capacity of detecting that a fault ϕ has occurred, a bounded *"time"* after its occurrence (dual properties concern the possibility to determine with certainty that ϕ has not occured).

[1] we only consider single fault types to avoid technicalities w.l.o.g.

III. Covering and Facets

Covering Relation: We recall here definitions and results from [14]. For a node $x \in (\mathcal{B} \cup \mathcal{E})$, the *conflict set* of x is $\#[x] \triangleq \{x' \mid x \# x'\}$. The *root conflict set* is given by

$$\underline{\#}[x] \quad \triangleq \quad \{y \mid x \# y \ \wedge \ \forall\, z : z < y \Rightarrow \neg\,(z \# x)\}\,.$$

As shown in [14], the set $\#[x]$ is generated by $\underline{\#}[x]$ through inheritance:

$$\#[x] \quad = \quad \left\{z \mid \exists\ y \in \underline{\#}[x] :\ y \leqslant z\right\}. \tag{4}$$

As a consequence, $x_1 \triangleright x_2$ iff $\underline{\#}[x_1] \supseteq \underline{\#}[x_2]$. Node x *leads*

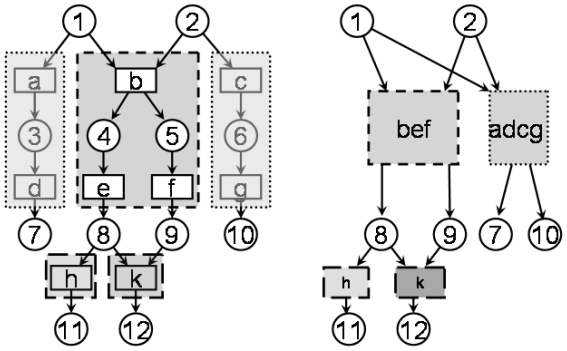

Fig. 2. *Left: Occurrence net with facets shown as rounded boxes; right: the occurrence net obtained from the left hand example by facet abstraction*

to or *covers* y, written $x \triangleright y$, iff $\#[x] \supseteq \#[y]$. Define the *covering range* of node x as

$$\triangleright[x] \quad \triangleq \quad \{y \mid x \triangleright y\};$$

we have (see [14]) the following properties of relation $\triangleright$:

1) $\triangleright$ is a reflexive and transitive relation.
2) $x \triangleright y$ holds iff for all runs ω,

$$x \in \omega \quad \Rightarrow \quad y \in \omega; \tag{5}$$

this motivates the expression *'a leads to b'* for $a \triangleright b$

3) $x < y$ implies that $y \triangleright x$.
4) $\triangleright[x]$ is a configuration.

In the occurrence net on the left hand side of Figure III, one has $a \triangleright d \triangleright a$, $d \triangleright c$, $e \triangleright f$, $f \not\triangleright h$, etc.

Facets: A **facet** of $\mathcal{ON}$ is a strongly connected component of $\triangleright$, i.e. a maximal set $\delta \subseteq (\mathcal{E} \cup \mathcal{B})$ of nodes such that for any $x, y \in \delta$, both $x \triangleright y$ and $y \triangleright x$ hold; we denote as $\delta(x)$ the unique facet that contains x, compare Figure III.We have (see [14]) that facets are convex, i.e. $x, y \in \delta$ and $x < y < z$ together imply $z \in \delta$; moreover, $y_1 \# y_2$ implies that $\delta_1 \neq \delta_2$, i.e. Facets are conflict-free. However, more is true: the quotient facet structure of $\mathcal{ON}$ is an occurrence net itself. In fact, let x_i be a node of $\mathcal{ON}$, $\delta_i \triangleq \delta(x_i)$, and

$$\delta_1 \prec_\Delta \delta_2 \quad \Longleftrightarrow \quad \left\{ \begin{array}{ll} & \delta_1 \neq \delta_2 \\ \exists & y_1 \in \delta_1, y_2 \in \delta_2 : \\ & y_1 < y_2 \end{array} \right. \tag{6}$$

$$\delta_1 \#_\Delta \delta_2 \quad \Longleftrightarrow \quad [\exists\ y_1 \in \delta_1, y_2 \in \delta_2 :\ y_1 \# y_2] \tag{7}$$

Relation $\prec_\Delta$ from Definition (6) is a partial order [14]. $\#_\Delta$ is well-defined since $y_1 \# y_2$ implies $z_1 \# z_2$ for all z_1 from δ_1 and z_2 from δ_2. One checks easily that

$$\delta_1 \# \delta_2 \prec_\Delta \delta_3 \quad \Rightarrow \quad \delta_1 \# \delta_3, \tag{8}$$

and finds that $(\Delta, \prec_\Delta, \#_\Delta)$ is an event structure in the sense of Definition 2. We denote as $\lceil \delta \rceil$ the set of facets

$$\lceil \delta \rceil \quad \triangleq \quad \{\delta' \mid \delta' \prec_\Delta \delta\}.$$

It can be shown (see [14]) that the set union of all facets in $\lceil \delta \rceil$ spans a configuration of $\mathcal{ON}$; we denote this configuration as

$$\kappa(\delta). \tag{9}$$

The Occurrence Net of Facets: Lumping facets into single events yields a new occurrence net $\mathfrak{F}(\mathcal{ON})$ [14]: Let $\mathcal{ON} = (\mathcal{B}, \mathcal{E}, G, \mathbf{c}_0)$ be an occurrence net, and Δ its set of facets. Set

$$\begin{array}{rcl} \mathfrak{F}(\mathcal{E}) & \triangleq & \Delta,\ \mathfrak{F}(\mathcal{B}) \triangleq \mathbf{c}_0 \cup \{b \mid b^\bullet \cap \delta(b) = \emptyset\} \\ \mathfrak{F}(G) & \triangleq & \{(b, \delta) \in \mathfrak{F}(\mathcal{B}) \times \mathfrak{F}(\mathcal{E}) \mid b^\bullet \cap \delta(b) \neq \emptyset\} \\ & & \cup\, \{(\delta, b) \in \Delta \times \mathfrak{F}(\mathcal{B}) \mid {}^\bullet b \subseteq \delta\}\,; \end{array}$$

then $\mathfrak{F}(\mathcal{ON}) = (\mathfrak{F}(\mathcal{B}), \mathfrak{F}(\mathcal{E}), \mathfrak{F}(G), \mathbf{c}_0)$ is an occurrence net.

Finitude of Facets: Note that in general, facets may be of infinite size; this motivates the following definition:

Definition 6: $\mathcal{N}$ is *universally non-deterministic (UND)* iff for any two distinct configurations κ, κ' of $\mathcal{N}$ such that κ is a prefix of κ' and $\kappa \ \sim_M \ \kappa'$, there must exist a configuration κ'' such that κ is a prefix of κ'', and for some $x' \in \kappa'$ and some $x'' \in \kappa''$ we have $x' \# x''$.

One checks that $\mathcal{U}_2(\mathcal{N})$ is sufficient to verify if $\mathcal{N}$ is UND. The important property of UND is:

Lemma 1: For UND nets, all facets are finite.

Proof: Suppose δ is infinite; then there exists at least one transition $t \in \mathcal{T}$ such δ contains an infinite chain $e_1 < e_2 < \ldots$ with $\forall\ i :\ \pi(e_i) = \mathcal{T}$. Let $M_i \triangleq M(\lceil e_i \rceil)$; since $\mathcal{N}$ is finite, it has only a finite number of reachable markings, hence there exists a reachable marking M and an infinite collection of indices $j_1 < j_2 < \ldots$ such that $\forall\ k :\ M_{j_k} = M$. If $\mathcal{N}$ is UND, there must therefore be a condition $b \in \delta \cap \mathcal{B}$ and events $e_a, e_b \in b^\bullet$ such that $e_{j_1} < b < e_a \leq e_{j_2}$, with $e_a \neq e_b$; therefore $e_a \not\triangleright e_{j_1}$, and hence $e_{j_2} \not\triangleright e_{j_1}$, contradicting the assumption that e_{j_2} and e_{j_1} are in the same facet. □

We return to the problem of diagnosability above, still keeping the same setting and notations. Let us define the *pro-cone* of a node $x \in \mathcal{E} \cup \mathcal{B}$ as

$$\lceil\lceil x \rceil\rceil \quad \triangleq \quad \kappa\left(\delta\left(x\right)\right); \tag{10}$$

the *closure* of a configuration κ is then

$$\lceil\lceil\kappa\rceil\rceil \triangleq \bigcup_{x\in\kappa'} \lceil\lceil x\rceil\rceil; \tag{11}$$

where we use the notation (9) for the configuration obtained from a cone of facets. Configuration κ is *closed* iff $\lceil\lceil\kappa\rceil\rceil = \kappa$. Notice that $\lceil\lceil\kappa\rceil\rceil$ coincides with the configuration obtained by intersecting all runs that extend κ; this makes closed configurations key entities for asynchronous diagnosis.

IV. $\mathbf{q}$-DIAGNOSABILITY

We are now ready for analyzing $\mathbf{q}-$diagnosability, restricted, for simplicity of presentation, to the case of one fault transition ϕ. First, let us correct the Definition from [14]:

Definition 7: In $\mathcal{N}$, a fault ξ is $\mathbf{q}$-*diagnosable* iff for any two configurations κ, κ' such that $\kappa \sim_M \kappa'$,

$$\lceil\lceil\kappa\rceil\rceil \sim_A \lceil\lceil\kappa'\rceil\rceil \;\Rightarrow\; \lceil\lceil\kappa\rceil\rceil \sim_\xi \lceil\lceil\kappa'\rceil\rceil. \tag{12}$$

In words, ξ is $\mathbf{q}$-diagnosable iff for any two configurations κ, κ' the following holds: if the *inevitable common parts*, as formalized by $\lceil\lceil\bullet\rceil\rceil$, of all runs extending κ and κ', respectively, are observationally equivalent, these closures have to be fault equivalent. Note that the above definition modifies the one given in [14], where the role of $\lceil\lceil\kappa'\rceil\rceil$ in formula (12) was played by κ' itself. That definition - while giving a notion that is interesting in its own right - is asymmetrical in the roles of κ and κ' and assumes a higher degree of observability than can be expected in a distributed system; we therefore recommend, and focus on, the above notion of $\mathbf{q}$-diagnosability. It is in fact well adapted to asynchronous systems: the precise interleaving of events is not available; the order of occurrence of concurrent events can not be observed, concurrent events will occur and go unnoticed *unless* they change future branchings. Therefore, the entity to inspect for alarms and faults are the *closed configurations*.

Consider the occurrence net of facets, $\mathcal{FN} \triangleq \mathfrak{F}(\mathcal{ON})$, where $\mathcal{ON}$ will be instantiated by the unfolding $\mathcal{U}(\mathcal{N})$ of the net $\mathcal{N}$ to be diagnosed. For a facet δ, let

$$\lceil\delta\rceil \in \mathbf{Con} \quad \triangleq \quad \mathbf{Con}(\mathcal{FN})$$

be the prime configuration of δ. Define the *immediate facet conflict relation* $\#_\mu^\Delta$ by:

$$\begin{array}{rll} \delta_1\#_\mu^\Delta\delta_2 \quad :\Longleftrightarrow & \exists\ \kappa & = \kappa(\delta_1,\delta_2) \in \mathbf{Con}: \\ & & \lceil\lceil\delta_1\rceil\rceil\backslash\kappa = \delta_1 \cup {\delta_1}^\bullet \\ & \wedge & \lceil\lceil\delta_2\rceil\rceil\backslash\kappa = \delta_2 \cup {\delta_2}^\bullet \\ & \wedge & \kappa_1 \triangleq \lceil\lceil\kappa\cup\delta_1\rceil\rceil \in \mathbf{Con} \\ & \wedge & \kappa_2 \triangleq \lceil\lceil\kappa\cup\delta_2\rceil\rceil \in \mathbf{Con} \\ & \wedge & \lceil\lceil\kappa\cup\{\delta_1,\delta\}\rceil\rceil \notin \mathbf{Con} \end{array}$$

Finally, call a facet δ ϕ-*negative* if δ contains no ϕ-event, and ϕ-*positive* otherwise.

With these preparations, we are now ready to define the properties that will characterize $\mathbf{q}$-Diagnosability. Let δ_A the partial order that facet δ induces on the set of its labeled events. $\mathcal{N}$ is *witnessful* iff every ϕ-positive facet δ contains a visible event $e \in \mathcal{E}_A$. We call $\mathcal{N}$ ϕ-*faithful* iff $\delta \sim_A \delta'$ implies that δ is ϕ-positive iff δ' is. $\mathcal{N}$ is *weakly If δ is ϕ-positive, then for any facet δ' such that $\delta\#_\mu^\Delta\delta'$ and δ_A is isomorphic to δ'_A, then δ' is also ϕ-positive.*

To state our main result, we say that the labeling λ of $\mathcal{N}$ is $\Delta-$*simple* iff every facet contains at most one observable event, i.e. from $\mathcal{E}_A$.

Theorem 4.1: If $\mathcal{N}$ is $\mathbf{q}$-diagnosable, then it is witnessful and weakly ϕ-faithful. Conversely, if λ is Δ-simple, witnessful and ϕ-faithful, then it is $\mathbf{q}$-diagnosable.

Proof: Suppose that $\mathcal{N}$ is $\mathbf{q}-diagnosable$ and has two facets δ, δ' with *(i)* δ_A and δ'_A isomorphic, and *(ii)* δ faulty; choose δ such that its past does not contain any copy of δ. As a consequence, δ also contains a first occurrence of ϕ. Let $\kappa(\delta,\delta')$ be a configuration as in the definition of $\#_\mu^\Delta$; then

$$\begin{array}{rcl} \kappa & \triangleq & \lceil\lceil\kappa(\delta,\delta')\cup\delta\rceil\rceil \\ \kappa' & \triangleq & \lceil\lceil\kappa(\delta,\delta')\cup\delta'\rceil\rceil \end{array}$$

satisfy $\lceil\lceil\kappa\rceil\rceil \sim_A \lceil\lceil\kappa'\rceil\rceil$ by construction. Hence, by $\mathbf{q}-$diagnosability, we must have $\lceil\lceil\kappa\rceil\rceil \sim_\phi \lceil\lceil\kappa'\rceil\rceil$, which implies, given our assumptions, that δ' must be faulty.
Now assume $\mathcal{N}$ is Δ-simple, witnessful and ϕ-faithful, and let κ be a faulty configuration, κ' a fault-free one, and $\lceil\lceil\kappa\rceil\rceil \sim_A \lceil\lceil\kappa'\rceil\rceil$. We assume w.l.o.g. that κ and κ' are both closed, and chosen minimal with these properties, i.e. that there is no closed configuration strictly smaller than κ for which the above holds for any choice of κ', and κ' is minimal for κ. This implies that there is a unique faulty facet δ such that $\kappa = \lceil\lceil\delta\rceil\rceil$. Since λ is simple, the first part of $\mathcal{D}$ yields that δ contributes exactly one letter to κ_A; call this letter a_δ. Let $\delta^1,\ldots,\delta^n$ be the maximal facets of κ'. Then, by minimality of κ', there is one index i such that δ^i_A is a_δ; thus the second part of $\mathcal{D}$ implies that δ^i is faulty, contradicting the assumption that κ' is fault-free. □

Theorem 4.1 allows to detect a vast class of diagnosability problems, and to ensure $\mathbf{q}$-diagnosability with fairly low effort on observability of events. Having one observable event in every faulty facet appears as both a natural and reasonable assumption for most application contexts such as physical plants: occurrence of a fault ϕ should be accompanied by an observable alarm, which may be emitted an unspecified time *later* than ϕ, but *independent* from any subsequent nondeterministic choice in the system.

Note that property UND is important to ensure identification of facets: thanks to Lemma 1, UND-nets have all facets finite, a property that is not satisfied in general.

The facet-wise diagnosis implied by the above results blends the traditionnaly disjoint tasks of *diagnosis* and *prognosis* in the sense that a ϕ-event may occur late in some facet δ, while the observable events allowing to recognize that δ is ϕ-positive may have occured long before.

In cases where the assumptions of Theorem 4.1 are 'oversatisfied' in the sense that more than one observation per

facet is available, it is often possible to deliberately 'forget' the excess information and reduce the problem to a Δ-simple case. Exploring this reduction is a topic of future work.

1) Performing **q**-*Diagnosis:* Qualitative *Diagnosis* can be performed for general UND-nets using diagnosers constructed from a suitable prefix of $\mathfrak{F}(\mathcal{U}_{\mathcal{N}})$, cut off and looped after repetition of markings. Consider the following equivalence relation on Δ: $\delta_1 \sim_\Delta \delta_2$ iff

1) δ_1 and δ_2 are isomorphic as π-labeled occurrence nets,
2) $\lceil \delta_1 \rceil \sim_M \lceil \delta_2 \rceil$, and
3) $\delta_1 \sim_A \delta_2$.

Denote the $\sim_\Delta$-equivalence class of $\delta \in \Delta$ as $\underline{\delta}$.

Setting, for $b_1, b_2 \in \mathfrak{F}(\mathcal{B})$, $b_1 \sim_\Delta b_2$ iff $\delta_1 \sim_\Delta$ for all facets δ_1, δ_2 such that $b_i \in {}^\bullet\delta_i$, $i \in \{1,2\}$; set $\mathcal{T}_{/_\Delta} \triangleq \Delta_{/\sim_\Delta}$ and $\mathcal{P}_{/_\Delta} \triangleq \mathfrak{F}(\mathcal{B})_{/\sim_\Delta}$, and $F_{/_\Delta} \triangleq (\mathfrak{F}(_)G)_{\sim_\Delta}$. From the construction of the unfolding $\mathcal{U}_{\mathcal{N}}$, we obtain that

1) $\mathcal{ON}_{/_\Delta} \triangleq (\mathcal{B}_{/_\Delta}, \mathcal{E}_{/_\Delta}, F_{/_\Delta}, M_{0/_\Delta})$ is a safe Petri net,
2) the unfolding of $\mathcal{ON}_{/_\Delta}$ is isomorphic (as a labelled occurrence net) to $\mathcal{ON}$.

The looped net thus obtained can be seen as a finite, facet-contracted version of $\mathcal{N}$. It allows to carry out diagnosability analysis, and diagnosis in the cases covered by Theorem 4.1, where the diagnosis procedure to recognition of local alarm patterns; we omit the formal diagnoser construction for lack of space here.

It is natural to ask whether, and how, diagnosis can be performed in cases where ϕ-positive facets can be silent (i.e. δ_A is empty). Theorem 4.1 gives a clear and negative answer to this: invisible facets destroy **q**-diagnosability. However, the above quotient construction allows to detect faulty invisible facets by diagnosis on the Petri net $\mathcal{ON}_{/_\Delta}$. In fact, the facets of $\mathcal{ON}$ becoming simple transitions of $\mathcal{ON}_{/_\Delta}$, one obtains a new Petri net diagnosis problem with a labelling $\overline{\lambda} : \Delta \to \mathfrak{A}$, where $\mathfrak{A}$ is the alphabet formed by the partially orderd shapes δ_A obtained under λ. In this quotient problem, it may be adequate to check **q**-diagnosability again, or e.g. weak diagnosability in the sense of [13]; this is a potential topic of interest in future work.

V. Conclusion and Outlook

q-diagnosability is an adequate property to require of asynchrounous systems in which state and time are distributed, and ordering of concurrent events is neither controllable nor observable. Facets formalize the fact that occurrence of some events reveals the inevitable occurrence of other events; the facet of an observable event e gives all the behaviours that will inevitably occur, sooner or later, if e does. One can effectively compute the basic relation $\triangleright$ verify $\mathbf{q}-$diagnosability using a finite unfolding prefix of bounded size, as our results here and in [14] show.

References

[1] J. Desel and J. Esparza. *Free Choice Petri Nets.* Cambridge University Press, 1995.

[2] A.T. Bouloutas, S. Calo, and A. Finkel. Alarm correlation and fault identification in communication networks. *IEEE Trans. on Communication* **42**(2-4), 1994.

[3] C. Cassandras and S. Lafortune. *Introduction to discrete event systems.* Kluwer Academic Publishers, 1999.

[4] V. Diekert and G. Rozenberg, eds. *The Book of Traces* . World Scientific, 1995.

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

[6] J. Esparza, S. Römer, W. Vogler. An improvement of McMillan's unfolding algorithm. *Formal Meth. in System Design* **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. *INRIA Research report* **5916**, Feb. 2007; http://hal.inria.fr/inria-00130025. Extended version of a plenary address at WODES 2006.

[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)·714–727, May 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, Mar. 2005

[10] R.G. Gardner and D. Harle. Methods and systems for alarm correlation. *GlobeCom 1996*.

[11] S. Genc and S. Lafortune. Distributed Diagnosis of Place-Bordered Petri Nets". *IEEE Transactions on Automation Science and Engineering*. April 2007

[12] M. P. Cabasino, A. Giua, C. Seatzu. Identification of Petri Nets from Knowledge of Their Language. Discrete Event Dynamic Systems **17**(4): 447–474, 2007.

[13] S. Haar, A. Benveniste, E. Fabre, and C. Jard. Partial Order Diagnosability of Discrete Event Systems Using Petri Net Unfoldings. In: *Proc. 42nd CDC,* 2003.

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

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

[16] C.N. Hadjicostis, and G.C. Verghese. Monitoring discrete event systems using Petri net embeddings. in *Proc. 20st (ICATPN)*, LNCS **1639**:188–208, Springer Verlag 1999.

[17] L.E. Holloway, B.H. Krogh and A. Giua. A Survey of Petri Net Methods for Controlled Discrete event systems. *Discrete Event Dynamic Systems: Theory and Applications* **7**:151–190, 1997.

[18] V. Khomenko, M. Koutny, and W. Vogler. Canonical Prefixes of Petri Net Unfoldings. *Acta Informatica* **40**:95–118, 2003. Preliminary Version in: D. Brinskma and K.G. Larsen (eds.), *Proc. CAV 2002*, LNCS **2404**:582–595. , Springer Verlag2002.

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

[20] K. McMillan. Using Unfoldings to avoid the state explosion problem in the verification of asynchronous circuits. *4th Workshop on Computer Aided Verification* 164–174, 1992.

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

[22] J.L. Peterson. *Petri Net Theory and the Modeling of Systems.* Prentice-Hall, 1981.

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

[24] A. Sahraoui, H. Atabakhche, M. Courvoisier, and R. Valette. Joining Petri nets and knowledge-based systems for monitoring purposes. *Proc. IEEE Int. Conf. on Robotics Automation,* 1160–1165, 1987.

[25] 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.

[26] R. Sengupta. Diagnosis and communications in distributed systems. *Proceedings* WODES *1998*, 144-151.

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

[28] 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.