Reveal Your Faults: It's Only Fair!

Stefan Haar, César Rodríguez, and Stefan Schwoon
INRIA and LSV, CNRS and ENS Cachan
61, av. du Président Wilson, 94235 Cachan Cedex, France
stefan.haar@inria.fr, cesar.rodriguez@inria.fr, stefan.schwoon@inria.fr

***Abstract*—We present a methodology for fault diagnosis in concurrent, partially observable systems with additional fairness constraints. In this *weak* diagnosis, one asks whether a concurrent chronicle of observed events allows to determine that a non-observable fault will inevitably occur, sooner or later, on any maximal system run compatible with the observation. The approach builds on strengths and techniques of unfoldings of safe Petri nets, striving to compute a compact prefix of the unfolding that carries sufficient information for the diagnosis algorithm. Our work extends and generalizes the unfolding-based diagnosis approaches by Benveniste et al. [1] as well as Esparza and Kern [2]. Both of these focused mostly on the use of sequential observations, in particular did not exploit the capacity of unfoldings to reveal inevitable occurrences of concurrent or future events studied by Balaguer et al. [3]. Our diagnosis method captures such indirect, revealed dependencies. We develop theoretical foundations and an algorithmic solution to the diagnosis problem, and present a SAT solving method for practical diagnosis with our approach.**

I. 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 this paper, we extend diagnosis to include detection of indirectly implied faults in concurrent systems, under an asynchronous, partial-order-based observation approach. Let us first review some settings for diagnosis.

The classical FSM setting [4]: Here one assumes that the observed system is an automaton or *finite state machine (FSM)* with transition set T, behavior given by a prefix-closed language $\mathcal{L} \subseteq T^*$, and a set of observable transition labels $\mathbb{O}$. The associated labeling $\lambda\colon T \to \mathbb{O}$ is not necessarily injective, and leaves some transitions from T unobservable, in particular a *fault* action ϕ[1]. Observations are words $O \in \mathbb{O}^*$, obtained by applying the 'mask' λ to words in T^*. Diagnosis is then the task of deciding whether or not all possible behaviors $w \in \mathcal{L}$ that *explain* an observation O, i.e., such that $w \in \lambda^{-1}(O)$, contain an occurrence of ϕ.

Concurrency and Asynchrony: Concurrent systems, e.g., in telecommunications, are difficult to supervise using the classical approach because of the state-explosion problem in using FSM models. Models that reflect the local and distributed nature of the observed system, such as Petri nets, are helpful not only in terms of computational efficiency, but also conceptually. Putting these ideas together, in [1] diagnosis is extended to asynchronous models and their non-interleaved semantics.[2] In *centralized non-sequential* (or asynchronous) *diagnosis*, there are several sensors, each of which observes (a fragment of) the system sequentially, as above. However, the streams of observations from the sensors reach the (centralizing) supervisor asynchronously; no assumption is made about the communication architecture or speed. One assumes that the architecture respects causality (if occurrence of a causally precedes that of a', the supervisor sees a before a'), and that the ordering of observations from the same sensor is respected.

This scenario must not be confused with either *(a) decentralized* or *(b) distributed diagnosis*, where several supervisors cooperate to reach a global verdict on whether or not a fault has occurred. In *(a)*, all supervisors emit local verdicts (e.g. yes/no) that are synthesized into a global one. In *(b)* [6], [7], *explanations* are computed in a distributed unfolding procedure, where supervisors build local explanations that are successively reduced by communication with other supervisors.

Here, we focus on extending *asynchronous diagnosis* of safe Petri nets with an unfolding-based approach that extends that of [1], [2] or [8]. The key idea is that in concurrent systems, certain events and properties can be 'implied', or *revealed*, by others that need not be in direct causal relation. The additional assumption we need is that of *weakly fair behavior*: on weakly fair runs, a transition t that becomes enabled at some point cannot stay enabled forever; eventually, either t or another conflicting transition t' must fire.[3] Call t', or t itself, a *spoiler* of t; for the run to be weakly fair, some spoiler of t must fire. We say that an observation O *(weakly) diagnoses* fault ϕ iff all weakly fair runs that explain O contain ϕ; this is in line with *weak diagnosability* from [9], [10], [11].

Both [1] and [2] use Petri net unfoldings under certain restrictions: [1] accepts partial-order observations, but refuses models with unobservable loops; [2] accepts the latter, thanks to dedicated *cutoff criteria*, but refuses the former. Our work uses both features, additionally accounting for weak fairness in the diagnosis procedure. We generalize the cutoff criteria of [2] to partially ordered observations, and extend them so as to guarantee in the constructed prefix a sufficient collection of *spoilers*, in the above sense.

Sec. II establishes basic notions and Sec. III presents our diagnosis framework. In Sec. IV, we develop theoretical foundations for the problem we solve. An algorithmic solution via a SAT-encoding is given in Sec. V and Sec. VI concludes and discusses future work.

[1] Or a *set* Φ of faults, which does not alter the problem much.
[2] Compare also the discussion of partial-order methods in [5].
[3] This is often not called 'weak fairness' but 'progress'.

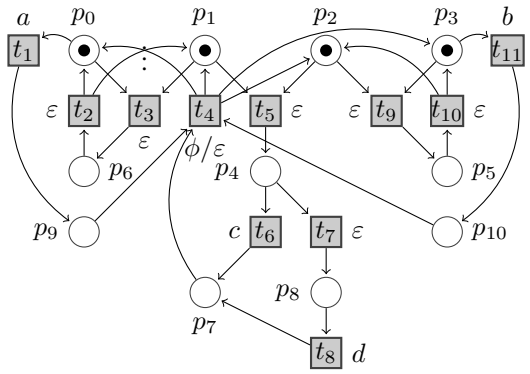

Fig. 1. A safe Petri net with partial observation. The inscription of a transition indicates its name; the label next to it is either a latin letter or empty (ε), in which case the transition is invisible. Transition t_4 is the invisible fault transition, which is called ϕ in the text and whose label is ε.

II. Basic notions

In this section, we establish notations and recall existing results about Petri nets and their unfoldings, see, e.g., [12]. We also discuss the notions of weak fairness and labelled partial orders, and prove some useful statements about them.

A. Petri nets

A *Petri net* (or simply *net*) is a tuple $N = \langle P, T, F, m_0 \rangle$, where P and T are disjoint sets of *places* and *transitions*, $F \subseteq (P \times T) \cup (T \times P)$ is the *flow relation*, and $m_0 \colon P \to \mathbb{N}$ is the *initial marking*. N is called *finite* if P and T are finite. Places and transitions together are called *nodes*. For $x \in P \cup T$, let ${}^\bullet x := \{y \in P \cup T \colon (y, x) \in F\}$ be the *preset*, and $x^\bullet := \{y \in P \cup T \colon (x, y) \in F\}$ the *postset* of x.

A *marking* of N is a function $m \colon P \to \mathbb{N}$ that assigns *tokens* to every place. A transition t is *enabled* at m if $m(p) \geq 1$ for all $p \in {}^\bullet t$. Such t can *fire*, leading to the new marking m' where $m'(p) = m(p) - |\{p\} \cap {}^\bullet t| + |\{p\} \cap t^\bullet|$; this is denoted as $m \xrightarrow{t} m'$. A finite sequence $\sigma = t_1 \dots t_n \in T^*$ is a *run* of N if there exist markings $m_1, \dots, m_n$ such that $m_0 \xrightarrow{t_1} m_1 \dots \xrightarrow{t_n} m_n$; marking m_n is the marking *reached by* σ, denoted by $mark(\sigma)$. A marking m is *reachable* if $m = mark(\sigma)$ for some run σ. Let $\mathbf{R}(N)$ be the set of reachable markings. An infinite sequence $t_1 t_2 \dots \in T^\infty$ is a run if all its finite prefixes are. N is *safe* if $m(p) \leq 1$ for all reachable m and $p \in P$. All Petri nets considered in this paper are safe, and we treat their markings as sets.

Fig. 1 shows a Petri net; as usual, places are drawn as circles, transitions as squares, and the number of tokens in every place represents the initial marking. Let $N' = \langle P', T', F', m'_0 \rangle$ be a Petri net. A *homomorphism* from N to N' is a function $h \colon P \cup T \to P' \cup T'$ satisfying $h(P) \subseteq P'$, $h(T) \subseteq T'$, $h(m_0) = m'_0$, and for all $t \in T$, h restricted to ${}^\bullet t$ and $t^\bullet$ is a bijection to ${}^\bullet h(t)$ and $h(t)^\bullet$, respectively.

B. Occurrence nets

The *causality* relation $<$ on a net N is the transitive closure of F and $\leq$ the reflexive closure of $<$. For a node x, we define its set of *causes* as $[x] := \{t \in T \colon t \leq x\}$. A set $X \subseteq T$ is *causally closed* if $[t] \subseteq X$ for all $t \in X$. The *conflict* relation $\# \subseteq (P \cup T)^2$ on N is the least symmetric relation satisfying

- $t \# t'$ if $t, t' \in T$ with $t \neq t'$ and ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$; and
- $x \# z$ if there is $y \in P \cup T$ such that $x \# y$ and $y < z$.

Two nodes x, y are *concurrent*, written $x \parallel y$, if neither $x \leq y$, nor $y \leq x$, nor $x \# y$ holds. A set $X \subseteq P \cup T$ is *conflict-free* if $\neg(x \# y)$ holds for all $x, y \in X$. An *occurrence net* is a safe net $O = \langle B, E, G, \tilde{m}_0 \rangle$ if $<$ is a strict partial order for O; for any $b \in B$, we have $|{}^\bullet b| \leq 1$; for all $e \in E$, $[e]$ is finite and $\neg(e \# e')$ holds; and $\tilde{m}_0 = \{\, b \in B \colon {}^\bullet b = \emptyset \,\}$.

As per tradition, we call the nodes of B *conditions*, and those of E *events*. A *configuration* of O is a causally closed, conflict-free set C of events. $\mathcal{C}(O)$ denotes the set of all such configurations. The configurations of O are *concurrent runs* of O, i.e., for every (finite or infinite) configuration $C = \{e_1, e_2, \dots\}$ of O there is at least one run $e_{i_1} e_{i_2} \dots$ of O, called *interleaving*, such that $e_{i_u} < e_{i_v}$ implies $u < v$. If C is finite, then all its interleavings reach the same marking $cut(C) := (\tilde{m}_0 \cup C^\bullet) \setminus {}^\bullet C$ of O. For $e \in E$, we say that C *enables* e, written $C \overset{e}{\rightsquigarrow}$, iff $e \notin C$ and $(C \cup \{e\}) \in \mathcal{C}(O)$.

The *height* of event e is recursively defined by $\mathcal{H}(e) := 1 + \max\{\mathcal{H}(e') \colon e' < e\}$, where $\max \emptyset := 0$. The height of a configuration C and of occurrence net O, provided both are finite, are defined as $\mathcal{H}(C) := \max\{\mathcal{H}(e) \colon e \in C\}$ and $\mathcal{H}(O) := \max\{\mathcal{H}(e) \colon e \in E\}$, respectively.

A *prefix* of O is a net $\mathcal{P} = \langle B', E', G', \tilde{m}_0 \rangle$ such that $E' \subseteq E$ is causally closed, $B' = \tilde{m}_0 \cup (E')^\bullet$, and G' is the restriction of G to $(B' \cup E')$; we denote this by $\mathcal{P} \sqsubseteq O$.

C. Unfoldings

A *branching process* of N is a pair $\mathcal{P} = \langle O, f \rangle$, where O is an occurrence net and f a homomorphism from O to N such that for all events e, e' of O, if ${}^\bullet e = {}^\bullet e'$ and $f(e) = f(e')$, then $e = e'$. For every N, there is a unique (up to isomorphism) maximal (w.r.t. $\sqsubseteq$) branching process $\mathcal{U}_N = \langle U, f' \rangle$ that we call the *unfolding* of N [12]. Thus, any branching processes $\langle O, f \rangle$ is characterised by a prefix O of U and the restriction f of f' to the elements of O. For convenience, we shall often equate a branching process with its underlying occurrence net and call it *unfolding prefix*. With finite configurations C of $\mathcal{U}_N$ we associate a marking of N, denoted $mark(C) := f(cut(C))$.

An unfolding prefix $\mathcal{P}$ is *marking-complete* if, any marking m is reachable in N iff there is $C \in \mathcal{C}(\mathcal{P})$ such that $m = mark(C)$. E.g., $\mathcal{U}_N$ is marking-complete but in general infinite.

If two finite configurations C, C' of $\mathcal{U}_N$ reach the same marking $mark(C) = mark(C')$, then the fragments of $\mathcal{U}_N$ '*coming after*' C and C' are isomorphic, since they are the unfoldings of N with the initial marking $mark(C)$. This is formalized, e.g., in Proposition 4.3 of [12]. Let h be such isomorphism. For every *extension* $I \subseteq E$ of C (i.e., $C \cap I = \emptyset$ and $C \cup I$ is a configuration of $\mathcal{U}_N$) there is an isomorphic

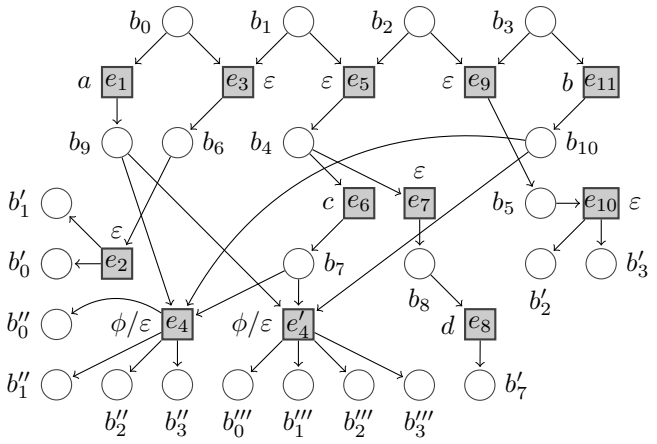

Fig. 2. A prefix of the unfolding for the net in Fig. 1. Events are named according to the corresponding transition (via the mapping f): e_i, e'_i, e''_i etc. for the occurrences of t_i. Observation labels are as in Fig. 1.

extension $I' := h(I)$ of C' reaching the same marking, i.e., $C' \cap I' = \emptyset$, and and $mark(C \cup I) = mark(C' \cup I')$.

D. Weakly fair runs and configurations

Runs represent a sequential view of the executions of a net, whereas configurations represent the concurrent point of view. We give and relate the definitions of weak fairness for both.

Let $N = \langle P, T, F, m_0 \rangle$ be a finite Petri net, and $\mathcal{U}_N = \langle \langle B, E, G, \tilde{m}_0 \rangle, f \rangle$ its unfolding. A configuration of $\mathcal{U}_N$ is *weakly fair* if it is $\subseteq$-maximal in $\mathcal{C}(\mathcal{U}_N)$. We denote by $\Omega(\mathcal{U}_N)$ the set of such configurations, or Ω if no confusion can arise.

Lemma 1. *A configuration ω is weakly fair iff it does not enable any event.*

Proof: If $\omega \in \Omega$ enables e, then $\omega \cup \{e\}$ is a configuration, and ω is not maximal, a contradiction. If ω is not weakly fair, then there exists a configuration C that is a proper superset of ω. Pick some $<$-minimal event e from $C \setminus \omega$. By assumption, all causal predecessors of e are in ω, so all conditions in ${}^\bullet e$ are either initial or receive a token from some event in ω. Since C is a configuration, it is conflict-free, so in particular no event in ω will remove a token from ${}^\bullet e$. Thus, ω enables e. ∎

Here we assume, for the sake of simplicity, that all weakly fair configurations of $\mathcal{U}_N$ are infinite. This entails no loss of generality, finite weakly fair configurations correspond to deadlocks that can be detected and processed separately — for instance, adding a looping transition.

A *spoiler* of transition t is any $t' \in T$ such that ${}^\bullet t \cap {}^\bullet t' \neq \emptyset$ (including t itself). We write $spoilers(t)$ for the set of such transitions. Following Vogler [13], who adapts the concept of *weakly fair termination* (see [14]) to a Petri net setting, we say that an infinite run $\sigma = t_1 t_2 \ldots \in T^\omega$ of N is *weakly fair* if its marking sequence $m_0, m_1, \ldots$ satisfies that for all $i \in \mathbb{N}$ and all $t \in T$, if m_i enables t, then there exists $j > i$ such that $t_j \in spoilers(t)$. In other words, any t enabled at some point along σ either fires eventually, or some other transition consumes from its preset. For runs of an occurrence net O, such as $\mathcal{U}_N$, we can make the following, stronger statement.

Lemma 2. *Let σ be a weakly fair run of O and $m_0, m_1 \ldots$ its marking sequence. For all $i \in \mathbb{N}$ and all $e \in E$, if m_i enables e, then $\exists k > i \; \forall j \geq k \colon m_j$ does not enable e.*

Proof: The statement follows from the definition of weakly fair runs and the fact that $<$ is acyclic for O; once a token from ${}^\bullet e$ is consumed, it cannot be replaced, and e remains disabled forever. ∎

Finally, we observe that weakly fair runs and weakly fair configurations are related in the following way:

Lemma 3. *Every weakly fair run of $\mathcal{U}_N$ is an interleaving of some $\omega \in \Omega$. Conversely, all interleavings of every $\omega \in \Omega$ are weakly fair.*

Proof: Let $\sigma = e_1 e_2 \ldots$ be an (infinite) weakly fair run and $\omega := \{ e_i \colon i \geq 1 \}$. Since no e_i can fire without its causal predecessors putting tokens into its preset, ω is causally closed. Due to the acyclic structure of $\mathcal{U}_N$, no condition can receive a token twice, no event will be repeated in σ, so no two events can consume from the same place, therefore ω is conflict-free and hence a configuration. Now suppose ω is not weakly fair, then by Lemma 1 it enables some event e. Arguing like in the proof of Lemma 1, we can conclude that all conditions in ${}^\bullet e$ are either initial or will receive tokens from events in ω, and that no event in ω consumes from ${}^\bullet e$. Thus, some marking m_i in the marking sequence for σ enables e, and then e is never disabled, which contradicts the weak fairness property for σ.

For the converse, let σ be an interleaving of a weakly fair configuration ω. Suppose that σ is not weakly fair; then some event e eventually becomes enabled during σ but neither e nor any conflicting event is in ω. Now, e can only become enabled if all its causal predecessors are in ω and put tokens into its preset, so $\omega \cup \{e\}$ is also causally closed. Thus, $\omega \cup \{e\}$ is a configuration, which contradicts maximality of ω. ∎

E. Labelled partial orders (LPOs)

An *alphabet* is a finite set $\mathbb{X}$, whose elements are called *letters*. A *labelled partial order* (LPO) over $\mathbb{X}$ is a tuple $\alpha = \langle S, <, \lambda \rangle$ where $< \subseteq S \times S$ is an irreflexive and transitive (hence antisymmetric) relation on S, and $\lambda \colon S \to \mathbb{X}$ a labelling map. The size $|\alpha|$ of α is $|S|$. Let $\alpha' = \langle S', <', \lambda' \rangle$ be an LPO over $\mathbb{X}$. A *homomorphism* from α to α' is a function $h \colon S \to S'$ verifying

- $\lambda(a) = \lambda'(h(a))$, and (1)
- $a < b$ implies $h(a) <' h(b)$ for all $a, b \in S$. (2)

An *isomorphism* between α and α' is a bijective homomorphism h from α to α' where h^{-1} is a homomorphism from α' to α. We say that α is *compatible* with α' if there exists a bijective function $f \colon S \to S'$ such that

- $\lambda(a) = \lambda'(f(a))$, and (3)
- $a < b$ implies $\neg(f(b) <' f(a))$ for all $a, b \in S$. (4)

Note: (2) and (4) are not equivalent, and f must be bijective. Denote by $compat(\alpha)$ the set of LPOs compatible with α.

Lemma 4. *Given LPOs α, α', if there is a bijective homomorphism h from α' to α, then $compat(\alpha) \subseteq compat(\alpha')$.*

Proof: Let $A := \langle S_A, <_A, \lambda_A\rangle \in compat(\alpha)$, and let $f_1 \colon S_A \to S$ be the associated bijection. Let $f_2 := h^{-1}$ be the (bijective) inverse of h. We show that $f := f_2 \circ f_1$ satisfies (3) and (4) from A to α'. For $a, b \in S_A$, we prove that:

- f is bijective, as it is the composition of two bijections.
- $\lambda_A(a) = \lambda'(f(a))$. This is because f_1 (by definition) and f_2 (by construction from h) preserve labels.
- $a <_A b$ implies $\neg(f(b) <' f(a))$. Assume that $a <_A b$ and $f(b) <' f(a)$. Let $a_\alpha := f_1(a)$, and $b_\alpha := f_1(b)$. We know that $\neg(b_\alpha < a_\alpha)$ holds by definition of f_1. But by definition of h also $b_\alpha = h(f(b)) < h(f(a)) = a_\alpha$ holds, a contradiction. ∎

Lemma 5. *LPOs α and α' are isomorphic iff $compat(\alpha) = compat(\alpha')$.*

Proof sketch: One direction is trivial, the other is reasoning by cases on the possible orderings in the LPOs, after establishing that S and S' have the same size. ∎

III. Reveals and Diagnosis

All diagnosis strives to detect '*hidden*' events, but we aim at diagnosing also *latent* but *inevitable* events, possibly in the future of the system's evolution. That is, we wish to diagnose exactly whether *every weakly fair run that is compatible with the observations thus far, contains a fault occurrence*. By the above, we thus need to consider all weakly fair configurations in Ω that contain an explanation of the current observation as a prefix. Let us formalize these notions.

A. Reveals Relations

In occurrence nets, for two events e, e', we say (see [15], [16], [9], [3], [17]) that e *reveals* e', written $e \triangleright e'$, iff $e \in \omega \Rightarrow e' \in \omega$ for all $\omega \in \Omega$, That is, the occurrence of e entails that e' *inevitably* will occur, or has already occurred. In Fig. 2, e.g., we have $e_5 \triangleright e_1$, $e_4 \triangleright e_6$, and $e_3 \triangleright e_2$. After the occurrence of e_5, the occurrence of e_3 has become impossible; in a weakly fair execution, e_1 must thus necessarily occur. Similarly, when e_4 occurs, e_6 must already have occurred previously ($\triangleright^{-1}$ includes and extends *causal precedence*). Also, all weakly fair configurations containing e_3 must contain e_2. This binary relation $\triangleright$ can be computed, for 1-safe Petri nets, on a finite prefix whose height is bounded [17].

Binary reveals helps in detecting invisible events; however, for diagnosis purposes, it is not strong enough. We shall need more general relation that relates sets to sets, namely the *extended-reveals* relation $\twoheadrightarrow$ introduced by Balaguer et al. [3]. Suppose that in Fig. 2, a and b are observable labels, and that we actually do observe their occurrence. From inspection of Fig. 2, (a) e_5 is bound to occur, and (b) one of the two instances of ϕ – e_4 or e'_4 – are inevitable. However, the binary relation $\triangleright$ does not permit to deduce (a) and (b): while the *conjunction* of e_1 and e_{11} makes e_5 inevitable, we have neither $e_1 \triangleright e_5$ nor $e_{11} \triangleright e_5$ individually; also, the *disjunction* of e_4 or e'_4 is certain once e_5 is assured, but neither e_4 or e'_4 are revealed individually. To account for such situations, following [3] we say, for sets of events A, B, that A *extended-reveals* B, written $A \twoheadrightarrow B$, iff every weakly fair configuration that contains A also '*hits*' B, i.e.,

$$A \twoheadrightarrow B \Leftrightarrow \forall \omega \in \Omega \colon \ (A \subseteq \omega \Rightarrow B \cap \omega \neq \emptyset)\,.$$

B. Diagnosis from Partial Observation

For the rest of the paper, we fix the following framework: Let $N = \langle P, T, F, m_0\rangle$ be a finite, safe Petri net with unfolding $\mathcal{U}_N = \langle\langle B, E, G, \tilde{m}_0\rangle, f\rangle$. We assume that all $t \in T$ have non-empty preset and that all $\omega \in \Omega$ are infinite. Let ε denote a unique '*empty*' symbol, and let $\mathbb{X}$ denote a non-empty *observation* (or *alarm*) *alphabet* where $\varepsilon \notin \mathbb{X}$. Let $\lambda \colon T \to \mathbb{X} \cup \{\varepsilon\}$ be the mapping associating transitions of N with observations or ε, and $\phi \in T$ the unique *fault* transition.

Let $T_{obs} := T \setminus T_{ubs}$ and $T_{ubs} := \lambda^{-1}(\varepsilon)$ be the *observable* and *unobservable transitions* of N. Furthermore, let $E_{obs} := f^{-1}(T_{obs})$ and $E_{ubs} := f^{-1}(T_{ubs})$ be the *observable* and *unobservable events* of $\mathcal{U}_N$. Naturally, we assume that ϕ is unobservable, i.e., $\phi \in T_{ubs}$, and define the set of *fault events* as $E_\phi := f^{-1}(\phi)$. Observe that $E_\phi \subseteq E_{ubs}$. We extend λ to E and, by abuse of notation, define $\lambda(e) := \lambda(f(e))$.

We now define the notion of *observation pattern*, or simply *observation*. Observations are LPOs over an observation alphabet. LPOs allow to capture linearly ordered observations, produced by a single sensor that observes the *interleaving* of the system; or sets of linear orders, produced by a number of sensors that locally observe each concurrent process of a distributed system; or yet others. This allows to work in an *asynchronous* setting as in [1], but without the need to enumerate the *interleavings* of observation patterns, not even in the correctness proofs, as opposed to the approach of [2]. For the rest of the paper, fix a *finite* observation pattern $\alpha := \langle S_\alpha, <_\alpha, \lambda_\alpha\rangle$ over the observation alphabet $\mathbb{X}$.

Given N and α, where α is the observation of some execution of N, our goal is to determine whether that execution contains a fault, assuming that runs of N are weakly fair. So we need to consider all those weakly fair configurations of $\mathcal{U}_N$ that are *compatible* with α. We formalize this in two steps. First, we associate every configuration C with an LPO $lpo(C) := \langle S, <', \lambda'\rangle$, where $S := C \cap E_{obs}$ are the observable events in C, $<'$ is the restriction of $\mathcal{U}_N$'s causal order $<$ to S, and $\lambda' \colon S \to \mathbb{X}$ is the restriction of λ to S. Since $<'$ and λ' are restrictions of $<$ and λ, it is safe to confuse them here, and so we will.

Second, we define the *observations* of C as the set $obs(C) := compat(lpo(C))$, i.e., the set of all (LPOs modelling) observations compatible to *the* LPO of C. Conversely, say that C *explains* observation α if $\alpha \in obs(C)$, and define

$$expl(\alpha) := \{C \in \mathcal{C}(\mathcal{U}_N) \colon \alpha \in obs(C)\}.$$

As a consequence of Lemma 5, for configurations C, C', we have $obs(C) = obs(C')$ iff $lpo(C)$ is isomorphic to $lpo(C')$.

Definition 1. *Observation pattern* α weakly diagnoses ϕ *iff*

$$\textit{for all } C \in \textit{expl}(\alpha),\ C \longrightarrow\!\!\!\!\triangleright\ E_\phi. \tag{5}$$

Since weak diagnosis is the only form of diagnosis we consider, we will simply speak of *diagnosis*.

In the context of Fig. 2, any observation containing a label c or d clearly diagnoses ϕ. What is more, and may serve to see the power of weak diagnosis here, is that observing a *and* b also weakly diagnoses ϕ: in fact, every configuration that allows to observe a and b must contain an occurrence of t_5 and the fault. For instance, observe that $\{e_1, e_{11}\} \longrightarrow\!\!\!\!\triangleright\ \{e_4, e'_4\}$ holds since every weakly fair run that contains e_1 and e_{11} must contain e_5 and thus either e_4 or e'_4. On the other hand, observing only a or even a sequence a^k is not sufficient for diagnosing ϕ (consider a weakly fair run composed only of occurrences of t_1, t_9 and t_{10}).

The *diagnosis problem* is to decide, given N and observation α, whether or not α weakly diagnoses ϕ. One of the keys to solve it is deciding relation (5) for a given configuration. We will show below that if C is finite, then $C \longrightarrow\!\!\!\!\triangleright\ E_\phi$ can be verified on a bounded extension of C. We give a formalized solution to this problem in Sec. IV.

IV. A SOLUTION USING EXTENDED REVEALS

By Def. 1 and definition of $\twoheadrightarrow$, α diagnoses ϕ iff

$$\forall C \in \textit{expl}(\alpha)\ \forall \omega \in \Omega\colon (C \subseteq \omega\ \Rightarrow\ E_\phi \cap \omega \neq \emptyset) \tag{6}$$

Swapping the two $\forall$, this can be equivalently rephrased as:

$$\forall \omega \in \Omega\colon\ (\exists C \in \textit{expl}(\alpha)\colon C \subseteq \omega)\ \Rightarrow\ E_\phi \cap \omega \neq \emptyset \tag{7}$$

In this section we derive an algorithm for deciding the *negation* of (7), i.e., given α, the algorithm decides whether or not there is some $\omega \in \Omega$ that contains an explanation $C \in \textit{expl}(\alpha)$ and such that $\omega \cap E_\phi = \emptyset$. Deriving this algorithm needs to overcome two obstacles:

1) $\textit{expl}(\alpha)$ may be infinite due to *unobservable loops*. In Fig. 2, a is explained by any configuration in which t_9 and t_{10} fire any number of times, followed by e_1.
2) Ω is an infinite set in general, which we need to finitely represent while still being able to check for set inclusion or whether $E_\phi \cap \omega \neq \emptyset$ for each weakly fair ω.

The main ideas behind our solution can be summarized as follows. Following [2], we fix 1) showing that it is *sufficient* for deciding (7) to search for C within a *finite subset* of $\textit{expl}(\alpha)$, instead of the entire, potentially-infinite set of explanations. Because this subset is finite, there exists an unfolding prefix that contains it entirely (P_α), we will see how to construct it. Once such C has been found, the algorithm needs to decide if it can be extended into a fault-free, weakly fair $\omega \in \Omega$. We show (Lemma 8) that this is the case iff two configurations $C_1 \subset C_2$ exist such that both of them reach the same marking, both are free of faults, C_2 disables every event enabled by C_1, and $C \subseteq C_1$. This result does not quite yet give an algorithm, as, e.g., C_2 could be unboundedly large. To fix this we define two *finite* unfolding prefixes $\mathcal{P}^1_N$ and $\mathcal{P}^2_N$ such that, first, $\mathcal{P}^1_N$ is contained in $\mathcal{P}^2_N$ and, second, the aforementioned C_1, C_2 exist iff configurations $\tilde{C}_1 \in \mathcal{C}(\mathcal{P}^1_N)$ and $\tilde{C}_2 \in \mathcal{C}(\mathcal{P}^2_N)$ exist and satisfy (again) that $\tilde{C}_1, \tilde{C}_2$ reach the same marking, both are free of faults, $\tilde{C}_2$ disables all events enabled by $\tilde{C}_1$, and some other technical condition asking (modulo details) that $C \subseteq \tilde{C}_1$. This 'iff' is shown in Lemma 10. These two prefixes can be seen as fixing 2). Our main result, Theorem 1, formalizes these ideas; Sec. V derives a decision procedure from them.

A. Succinct explanations

Following [2], we define a finite subclass of explanations of α and show that it is *sufficient* for deciding (7).

Definition 2. *Configuration* $C \in \mathcal{C}(\mathcal{U}_N)$ *is* verbose *if it contains two events* $e, e' \in C$ *satisfying*

- $e' < e$, *and* (8)
- $\mathit{mark}([e']) = \mathit{mark}([e])$, *and* (9)
- $\mathit{obs}([e']) = \mathit{obs}([e])$. (10)

If C *is not verbose, it is* succinct.

Intuitively, C is verbose if it contains the occurrence of some loop of N consisting entirely of unobservable transitions. In Fig. 1, t_2, t_3 is an unobservable loop, that produces the verbose configuration $\{e_1, e'_3, e'_2\}$ in Fig. 2 (set e, e' of Def. 2 as $e := e'_2$ and $e' := e_1$).

Observe that $\mathit{lpo}([e'])$ is isomorphic to $\mathit{lpo}([e])$, by (10) and Lemma 5, and that $[e'] \subseteq [e]$, by (8). This means that all events in $[e] \setminus [e']$ are unobservable, see Fig. 3 (a). It also means that $\mathit{lpo}([e]) = \mathit{lpo}([e'])$, i.e., the LPO isomorphism is the identity function on C restricted to observable events. Finally, observe that (9) does *not* imply (8), even for 1-safe nets.

Def. 2 is different from the equivalent definition in [2] only in that $\mathit{obs}(C)$ is defined differently: here, it is a set of LPOs while in [2] it is the set of sequences $\lambda(\sigma)$ where σ is an interleaving of C.

Every *finite* verbose configuration C can be *peeled* (eventually multiple times) yielding a succinct configuration that reaches the same marking and has at least the same observations. If C is an explanation of some α, peeling intuitively corresponds to finding a shorter explanation C' where unnecessary (unobservable) fragments of C have been removed. Although $C' \subseteq C$ may hold, in general it does *not* hold.

Let us formalise this idea. Let C be verbose and finite, and let $e, e' \in C$ be events satisfying (8)-(10). Let $I := C \setminus [e]$. Recall that I are events of $\mathcal{U}_N$ that fire after $\mathit{cut}([e])$. We define $\mathit{peel}_{e,e'}(C)$ as the configuration

$$C' := [e'] \cup I',$$

where I' is the isomorphic copy of I that *continues* $\mathcal{U}_N$ just after $\mathit{cut}([e'])$. Recall that I' is well defined, as we said in Sec. II-C. Since $|[e']| < |[e]|$, we have $|C'| < |C|$. So if C' is still verbose, we only need to peel again finitely many times before obtaining a succinct configuration. We denote by $\mathit{peel}^*(C)$ the set of succinct configurations resulting from peeling C as many times as necessary (choosing any e, e' that satisfy (8)-(10) every time we peel). We conjecture that $\mathit{peel}^*(C)$ is a singleton, but do not rely on it in the sequel. Lemma 6 implies that $\mathit{peel}^*(C)$ are explanations of α if C is.

Lemma 6. *For any verbose configuration C with $C' := peel_{e,e'}(C)$, it holds that:*

- $mark(C) = mark(C')$ *(11)*
- $obs(C) \subseteq obs(C')$ *(12)*
- $C' \cap E_\phi \neq \emptyset \;\Rightarrow\; C \cap E_\phi \neq \emptyset$ *(13)*

Proof: Let $e, e' \in C$ be events satisfying (8)-(10). Recall that C has the form $[e] \uplus I$, and C' the form $[e'] \uplus I'$.

(11) is a consequence of (9) and the fact that I, I' are isomorphic. (12) is more laborious. Let $lpo(C) := \langle S, <, \lambda\rangle$ and $lpo(C') := \langle S', <, \lambda\rangle$. In the sequel we define a mapping $h\colon S' \to S$ and prove that h is a bijective homomorphism from $lpo(C')$ to $lpo(C)$. (12) then follows by Lemma 4.

Let f_1 be the LPO isomorphism between $lpo([e])$ and $lpo([e'])$. Recall that f_1 is actually the identity function. Let $f_2\colon I' \to I$ be the isomorphism between I' and I. Define $h := f_1 \cup f_2'$ where f_2' is the restriction of f_2 to S', i.e., the observable events of C'.

Observe that h is bijective because it is the union of two bijections whose domains and codomains are disjoint; it satisfies (1) because so do f_1 and f_2. Finally, for $e_1, e_2 \in S'$ with $e_1 < e_2$, we show that (2) holds. There are three cases:

- $e_1, e_2 \in [e']$. Then $h(e_1) = e_1 < e_2 = h(e_2)$.
- $e_1, e_2 \in I'$. Since the isomorphism f_2 preserves causality, we have $h(e_1) = f_2(e_1) < f_2(e_2) = h(e_2)$.
- $e_1 \in [e']$ and $e_2 \in I'$. There is some $c \in cut([e'])$ such that $e_1 < c < e_2$. Since N is safe and by (9), there is a single condition $c' \in cut([e])$ such that $f(c) = f(c')$, where f is $\mathcal{U}_N$'s labelling. $c \# c'$ or $c \parallel c'$ does not hold, because $[c] \cup [c'] \subseteq C$ and N is safe. So $c \leq c'$ holds, since $c' < c$ contradicts $e' < e$. Since $[e_2]$ consumes c, necessarily $[f_2(e_2)]$ consumes c', as I' and I are isomorphic. We therefore have:

$$h(e_1) = f_1(e_1) = e_1 < c \leq c' < f_2(e_2) = h(e_2)$$

As for(13), let $\tilde{e} \in C' \cap E_\phi$ be some fault. If $\tilde{e} \in [e'] \subseteq C$, then $\tilde{e} \in C$. If $\tilde{e} \in I'$, then $f_2(\tilde{e}) \in I \subseteq C$ is also fault, because f_2 preserves transition labels. In both cases $C \cap E_\phi \neq \emptyset$. ∎

Define the set of succinct explanations of α as

$$succexpl(\alpha) := \{C \in expl(\alpha) \mid C \text{ is succinct}\}$$

Proposition 1. *$succexpl(\alpha)$ is finite.*

Proof: Since N is finite, there are finitely many events in $\mathcal{U}_N$ of height less or equal to any given $n \in \mathbb{N}$, and thus finitely many configurations made up of such events. Assume now there are infinitely many succinct explanations of α. Because of the above, there must be a succinct explanation C that contains $e \in C$ such that $\mathcal{H}(e) = k(|\alpha| + 1)$, where k is the number of reachable markings in the net. Let $e_1 < \ldots < e_{|\alpha|+1} = e$ be events in some causality chain from $\tilde{m}_0$ to e such that $mark(e_1) = \ldots = mark(e_{|\alpha|+1})$, which exist by the pigeonhole principle. Since only $|\alpha|$ events in C are observable, $[e_{i+1}]\setminus[e_i] \subseteq E_{ubs}$ holds for some i. So C contains two causally related events, whose local configurations have the same LPO, thus the same observation (Lemma 5), and reach the same marking. So C is verbose, a contradiction. ∎

The previous proof works even if (8) is removed from Def. 2. However, (8) is required for the following statement.

Proposition 2. *The shortest explanations of α are succinct.*

Proof: (Shortest in number of events) If C explains α and is verbose, $peel^*(C)$ are shorter explanations of α. ∎

Proposition 2 would still work if (8) is replaced by the more general condition $|[e']| < |[e]|$, but (12) would become false.

The main lemma of this subsection shows that (7) can be rephrased into (14), eliminating the need to examine all potentially infinitely many explanations of α.

Lemma 7. *Observation pattern α diagnoses ϕ iff*

$$\forall \omega \in \Omega\colon\ (\exists C \in succexpl(\alpha)\colon C \subseteq \omega) \Rightarrow E_\phi \cap \omega \neq \emptyset \quad (14)$$

Proof: (7) and (14) differ only in that $succexpl(\alpha)$ has replaced $expl(\alpha)$. Trivially, (7) implies (14). Assume that (14) holds and let $\omega \in \Omega$, $C \in expl(\alpha)$ be such that $C \subseteq \omega$. We show that $\omega \cap E_\phi \neq \emptyset$. If C is succinct, we are done, so assume C is verbose and let $C' \in peel^*(C)$. By (11), we can append an isomorphic copy I of $\omega \setminus C$ to C', yielding a weakly fair configuration $\omega' := C' \cup I$. By (12), C' is a succinct explanation of α. Because $C' \subseteq \omega'$ and (14), it holds that $\omega' \cap E_\phi \neq \emptyset$. If $C' \cap E_\phi \neq \emptyset$, by (13), we have $C \cap E_\phi \neq \emptyset$. If $I \cap E_\phi \neq \emptyset$, then $\omega \setminus C$ must contain a fault because it is isomorphic to I. In any case, $\omega \cap E_\phi \neq \emptyset$. ∎

B. Characterizing weakly fair configurations

We now investigate a finite characterization of the weakly fair configurations in Ω that allows to reason about (*i*) inclusion of (succinct) explanations and (*ii*) absence of faults.

According to (14), α does *not* diagnose ϕ iff we can find a fault-free, weakly fair configuration ω that contains a succinct explanation. The next lemma establishes a characterization of such configurations, where the spoilers of an event play an important role:

Lemma 8. *Let C be a finite configuration. There exists a weakly fair $\omega \in \Omega$ such that $C \subseteq \omega$ and $\omega \cap E_\phi = \emptyset$ iff there are $C_1, C_2 \in \mathcal{C}(\mathcal{U}_N)$ satisfying*

- $C \subseteq C_1 \subseteq C_2$, *and* *(15)*
- $mark(C_1) = mark(C_2)$, *and* *(16)*
- $\forall e \in E\colon C_1 \overset{e}{\rightsquigarrow} \;\Rightarrow\; spoilers(e) \cap C_2 \neq \emptyset$, *and* *(17)*
- $C_2 \cap E_\phi = \emptyset$. *(18)*

Proof: The main idea is simple: fault-free, weakly fair configurations exist iff one can find configurations C_1, C_2 with $C_1 \subseteq C_2$, both free of faults, extending C, reaching the same marking, and such that $C_2 \setminus C_1$ disables all events enabled by C_1. The fragment $C_2 \setminus C_1$ can then be iterated infinitely often without leaving any event enabled forever.

Formally, let $\omega \in \Omega$ be weakly fair, such that $C \subseteq \omega$ and $\omega \cap E_\phi = \emptyset$. Let $\sigma = e_1 e_2 \ldots$ be any weakly fair interleaving of ω, and e_n the last event of C in σ. By the pigeonhole principle there is infinitely many $n \leq n_1 < n_2 < \cdots \in \mathbb{N}$ such that $mark(\sigma_{n_1}) = mark(\sigma_{n_2}) = \ldots$, where σ_i denotes the run $e_1 e_2 \ldots e_i$. Let C_1 be the restriction of ω to σ_{n_1}. Because

σ is weakly fair, there is some $i \in \mathbb{N}$ such that σ_{n_i} contains one spoiler for every event enabled by C_1 (Lemma 2). Let C_2 be the restriction of ω to σ_{n_i}. Then C_1, C_2 satisfy (15)-(18).

For the opposite direction, let C_1, C_2 be configurations satisfying (15)-(18). We construct a fault-free, weakly fair $\omega \in \Omega$. For convenience, we write $\mathbf{PLoop}(C_1, C_2)$ for all pairs $C_1, C_2 \in \mathcal{C}$ satisfying (15) and (16). Since $mark(C_1) = mark(C_2)$, we can append an isomorphic copy of $C_2 \setminus C_1$ to C_2, yielding C_3, such that $\mathbf{PLoop}(C_2, C_3)$ and $C_3 \cap E_\phi = \emptyset$ hold. Iterating this construction, we obtain a family $(C_n)_{n\in\mathbb{N}}$ of configurations satisfying $\mathbf{PLoop}(C_n, C_{n+1})$ and $C_n \cap E_\phi = \emptyset$ for all $n \in \mathbb{N}$. Let $\omega := \bigcup_{n\in\mathbb{N}} C_n$ be the configuration resulting from their union. We prove that ω does not enable any event, and thus $\omega \in \Omega$ by Lemma 1. By contradiction, let e be any event enabled by ω. It is therefore enabled by C_i for some $i \in \mathbb{N}$. Since the unfolding after $cut(C_i)$ is isomorphic to the unfolding after $cut(C_1)$, there is some e' isomorphic to e that is enabled by C_1. By construction, $spoilers(e') \cap C_2 \neq \emptyset$, so there is some $\widehat{e'}$ in $C_2 \setminus C_1$ that disables e'. Because $C_{i+1} \setminus C_i$ is isomorphic to $C_2 \setminus C_1$, there is some spoiler of e in $C_{i+1} \setminus C_i$, and ω does not enable e, a contradiction. ■

In Fig. 1, $(t_1, t_9, t_{10})^\omega$ is a weakly fair run, represented in Fig. 2 by the fault-free, weakly fair configuration $\omega := \{e_1, e_9, e_{10}, e'_1, \ldots\} \in \Omega$. Setting $C := \emptyset$, Lemma 8 implies the existence of $C_1 = \emptyset$ and $C_2 = \{e_1, e_9, e_{10}\}$.

While Lemma 8 identifies a method for finding fault-free, weakly fair configurations, there are still infinitely many configurations C_1, C_2 to consider. We would thus like to define a *finite* unfolding prefix of $\mathcal{U}_N$ such that, C_1, C_2 exist and verify (15) to (18) iff there are *small copies* of C_1, C_2 among the configurations of such a prefix that still satisfy (15) to (18).

Recall that it is possible [12] to compute a finite, marking-complete prefix $\mathcal{P}^1_N := \langle B_1, E_1, G_1, \tilde{m}_0\rangle$ of $\mathcal{U}_N$. For any configuration C of $\mathcal{U}_N$, we denote by C_{E_1} the set $C \cap E_1$.

Below, we define a prefix $\mathcal{P}^2_N$ that includes $\mathcal{P}^1_N$ and preserves not only reachability of markings but also the capability of a configuration to *spoil* previously enabled events, i.e., those events that can take the role of C_2 in (17). $\mathcal{P}^2_N$ will be defined using the following notion of cutoff, which is relative to $\mathcal{P}^1_N$:

Definition 3. *Event $e \in E$ is an* sp-cutoff *if there is $e' \in E$ such that, setting $D := [e] \setminus [e']$, we have $e' < e$, and*

- $f(^\bullet D \setminus D^\bullet) = f(D^\bullet \setminus {}^\bullet D)$, *and* (19)
- $B_1 \cap {}^\bullet D = \emptyset$. (20)

Observe that (20) imposes that D, the 'difference' between $[e]$ and $[e']$, cannot consume conditions generated by events of $\mathcal{P}^1_N$, see Fig. 3 (b). This means that $[e]$ and $[e']$ spoil exactly the same events among all those enabled by any $C_1 \in \mathcal{C}(\mathcal{P}^1_N)$. Although it is not clear after Def. 3, (19) implies $mark([e]) = mark([e])$, as the following lemma implies:

Lemma 9. *Let e be an sp-cutoff, and e' as in Def. 3. For all $C \in \mathcal{C}(\mathcal{P}^1_N)$, if $C \cup [e]$ is a configuration, then*

$$mark(C \cup [e]) = mark(C \cup [e']).$$

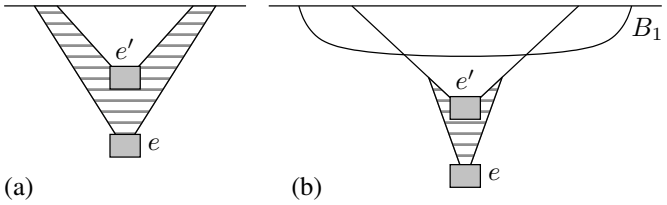

Fig. 3. The lined area in (a) represents unobservable, ϕ-labelled, events of a verbose explanation. In (b), the lined area depicts $D := [e] \setminus [e']$ in Def. 3, which does not consume any condition in B_1.

Proof: (sketch) Show by cases, using (19)-(20), that

$$f((\tilde{m}_0 \cup C^\bullet \cup [e]^\bullet) \setminus ({}^\bullet C \cup {}^\bullet[e])) = f((\tilde{m}_0 \cup C^\bullet \cup [e']^\bullet) \setminus ({}^\bullet C \cup {}^\bullet[e'])).$$ ■

Any configuration C of $\mathcal{U}_N$ that contains some sp-cutoff e can be *trimmed* in a way analogous to the way verbose configurations can be peeled into succinct configurations, cf. Lemma 6. Trimming C corresponds to finding some smaller configuration C' that preserves the above spoiling capabilities.

Formally, let C be any configuration that contains an sp-cutoff e and event $e' \in C$ as in Def. 3. Consider the configuration $C_{E_1} \cup [e] \subseteq C$. Since C_{E_1} is a configuration of $\mathcal{P}^1_N$, by Lemma 9, $mark(C_{E_1} \cup [e]) = mark(C_{E_1} \cup [e'])$. So we can partition C as $(C_{E_1} \cup [e])$ and $I := C \setminus (C_{E_1} \cup [e])$, and define $trim_{e,e'}(C)$ as the configuration

$$C' := (C_{E_1} \cup [e']) \cup I',$$

where I' is the isomorphic copy of I after $cut(C_{E_1} \cup [e'])$.

Lemma 10. *Let $C' := trim_{e,e'}(C)$ for any configuration C of $\mathcal{U}_N$. We have:*

- $mark(C) = mark(C')$ (21)
- $C \cap E_1 \subseteq C' \cap E_1$ (22)
- $C \cap E_\phi = \emptyset \Rightarrow C' \cap E_\phi = \emptyset$ (23)
- $\forall e \in B_1^\bullet$: $spoilers(e) \cap C \neq \emptyset \Rightarrow spoilers(e) \cap C' \neq \emptyset$ (24)
- $|C'| < |C|$ (25)

Proof: (22), (23), and (25) hold by construction of C'. (21) is a consequence of the fact that $\mathcal{U}_N$ stripped of $C_{E_1} \cup [e]$ is isomorphic to $\mathcal{U}_N$ stripped of $C_{E_1} \cup [e']$. So isomorphic sets of events I and I' yield the same marking of N. As for (24), let $e, e' \in C$ be as in Def. 3, let $\hat{e} \in B_1^\bullet$, and let $\hat{e}^\dagger \in spoilers(\hat{e}) \cap C$. Three cases are possible:

- $\hat{e}^\dagger \in (C \cap E_1) \cup [e']$. Then by construction, $\hat{e}^\dagger \in C'$.
- $\hat{e}^\dagger \in [e] \setminus [e']$. Not possible, entails contradiction to (20).
- $\hat{e}^\dagger \in I$. Let $c \in {}^\bullet\hat{e}^\dagger \cap B_1$ be any condition in B_1 consumed by $\hat{e}^\dagger$. We show that $c \in cut(C_{E_1} \cup [e'])$. This is because $c \in \tilde{m}_0 \cup (C \cap E_1)^\bullet$ (since C is causally closed), and also $c \notin {}^\bullet(C_{E_1} \cup [e'])$ (since the only event $\hat{e}^\dagger$ in C that consumes c is in I). With analogous reasoning, one shows that $c \in cut(C_{E_1} \cup [e])$. Now, because I, starting from $cut(C_{E_1} \cup [e])$, is isomorphic to I', starting from $cut(C_{E_1} \cup [e'])$, and c belongs to both cuts, if $\hat{e}^\dagger$ consumes from c, its isomorphic event in I' also consumes

from c and hence spoils $\hat{e}$. ∎

Observe that no event in $\mathcal{P}_N^1$ is an sp-cutoff, and thus $\mathcal{P}_N^2$ contains $\mathcal{P}_N^1$. Trimming decrements the number of events (25), so if $trim_{e,e'}(C)$ still has a sp-cutoff we can trim again finitely many times until getting an sp-cutoff-free configuration, choosing any e, e' as in Def. 3 every time we trim. Let $trim^*(C)$ denote the set of such configurations.

We define $\mathcal{P}_N^2 := \langle B_2, E_2, G_2, \tilde{m}_0\rangle$ as the unfolding prefix whose events are exactly all non sp-cutoff events, i.e.,

$$E_2 := \{e \in E \colon e \text{ is not sp-cutoff}\}.$$

Proposition 3. *$\mathcal{P}_N^2$ is finite.*

Proof: Assume E_2 is infinite. As in the proof of Proposition 1, we can find infinitely many events $e_1 < e_2 < \ldots$ in E_2. Because T and the number of reachable markings in N are finite, we can furthermore assume that $f(e_1) = f(e_2) = \ldots$ and $mark([e_1]) = mark([e_2]) = \ldots$ Define the sequence of *difference* sets $D_i = [e_i] \setminus [e_{i+1}]$, for $i \geq 1$. For $i < j$, ${}^\bullet D_i \cap {}^\bullet D_j = \emptyset$, otherwise $[e_j]$ would have conflicts. Since B_1 is finite, the number of D_is consuming from B_1 must then be finite. So there is some $k \geq 1$ such that $B_1 \cap {}^\bullet D_i = \emptyset$ holds for all $i \geq k$. Then e_{k+1} is an sp-cutoff, a contradiction. ∎

We can now state our main result:

Theorem 1. *Observation pattern α does* not *diagnose ϕ iff there exist configurations*

$$C, C_1' \in \mathcal{C}(\mathcal{U}_N), \qquad C_1 \in \mathcal{C}(\mathcal{P}_N^1), \qquad C_2 \in \mathcal{C}(\mathcal{P}_N^2)$$

satisfying the following properties:

- *C is a succinct explanation of α, and* (26)
- *$C \subseteq C_1'$, and* (27)
- *$C_1 \subseteq C_2$, and* (28)
- *$mark(C_1') = mark(C_1) = mark(C_2)$, and* (29)
- *$\forall e \in E \colon C_1 \overset{e}{\rightsquigarrow} \Rightarrow spoilers(e) \cap C_2 \neq \emptyset$, and* (30)
- *there is no fault event in either C_1' or C_2.* (31)

Proof: By (14), if α does not diagnose ϕ, there is a fault-free, weakly fair configuration $\omega \in \Omega$ and some succinct explanation $C \in succexpl(\alpha)$ with $C \subseteq \omega$. By Lemma 8, there are configurations $\tilde{C}_1, \tilde{C}_2$ that satisfy (15)-(18). Define $C_1 \in \mathcal{C}(\mathcal{P}_N^1)$ as any configuration in $\mathcal{P}_N^1$ that reaches $mark(\tilde{C}_1)$. Now let C_2' denote $C_1 \cup I$ where I is an isomorphic copy of $\tilde{C}_2 \setminus \tilde{C}_1$ starting at $cut(C_1)$. Define C_2 as either C_2' if C_2' contains no sp-cutoff or as any configuration in $trim^*(C_2')$ otherwise. In both cases, $C_2 \in \mathcal{C}(\mathcal{P}_N^2)$. Define $C_1' \subseteq \omega$ as any configuration satisfying (27) whose marking is $mark(\tilde{C}_1)$, which exists because ω repeats $mark(\tilde{C}_1)$ infinitely often.

(26) and (27) holds by definition of C, C_1'. By construction, $C_1 \subseteq C_2'$. If C_2' has sp-cutoffs and C_2 is taken from $trim^*(C_2')$, by (22) and the fact that $C_1 \subseteq E_1$, we have $C_1 \subseteq C_2$. So (28) holds in any case. (29) holds by construction of C_1', C_1, (16), and (21). Because ω is fault-free, C_1' is as well. By (18), $\tilde{C}_1, \tilde{C}_2$ are fault free, and so is C_2' (by isomorphism). Then, by (23), C_2 is fault-free. This shows (31). As for (30), we observe the following. $\tilde{C}_1, \tilde{C}_2$ satisfy (17). Since $\tilde{C}_1$ cannot contain any spoiler of the events it enables, all such spoilers are in $\tilde{C}_2 \setminus \tilde{C}_1$. Then, C_2' disables all events enabled by C_1 because $C_2' \setminus C_1$ is isomorphic to $\tilde{C}_2 \setminus \tilde{C}_1$. Now, because $C_1 \subseteq E_1$, all such events are in $B_1^\bullet$. Then by (24), C_2 disables all them, and (30) holds.

If C, C_1', C_1, C_2 exist and verify (26)-(31), by Lemma 8 some fault-free, weakly fair configuration $\omega \in \Omega$ exists and repeats infinitely often $mark(C_1) = mark(C_1')$. Construct another weakly fair configuration $\omega' := C_1' \cup I \in \Omega$ where I is an isomorphic copy of $\omega \setminus C_1$ starting at $cut(C_1')$. Now, ω' contains a succinct explanation and is fault-free because so was C_1', and by isomorphism between $\omega \setminus C_1$ and I. ∎

In other words, Theorem 1 states that α does not diagnose ϕ iff one can find configurations C_1, C_2 in some suitable *finite* unfolding prefixes and a succinct explanation C of α such that $mark(C_1)$ can be reached from $mark(C)$ without executing fault events. There is only finitely many succinct explanations of α, and we can decide whether one marking is reachable from another without executing faults using our next Proposition. So Theorem 1 suggests a decision algorithm for the diagnosis problem that we shall investigate in Sec. V.

Proposition 4. *There exist fault-free configurations $C \subseteq C'$ of $\mathcal{U}_N$ iff there is fault-free configurations $\hat{C} \subseteq \hat{C}'$ of, respectively, $\mathcal{P}_N^1, \mathcal{P}_N^2$, satisfying:*

$$mark(C) = mark(\hat{C}) \quad \text{and} \quad mark(C') = mark(\hat{C}').$$

Proof sketch: $\mathcal{P}_N^1$ is marking-complete and we can find the requested fault-free $\hat{C}$ in $\mathcal{C}(\mathcal{P}_N^1)$, see the proof of Proposition 4.9 (a) in [12]. Let $\hat{C}'' := \hat{C} \cup I$ where I is an isomorphic copy of $C' \setminus C$ starting at $cut(\hat{C})$. Let $\hat{C}' := \hat{C}''$ if $\hat{C}'' \subseteq E_2$, or $\hat{C}' \in trim^*(\hat{C}'')$ otherwise. Then $\hat{C}, \hat{C}'$ satisfy the lemma by (21)-(23). ∎

V. A DECISION METHOD FOR DIAGNOSIS

Theorem 1 states a set of necessary and sufficient conditions that characterize whether or not a given observation α diagnoses ϕ. In this section, we present a method for deciding if these conditions hold. We discuss (Sec. V-A) which information is needed in order to decide them, and how to obtain that information (Sec. V-B). Based on this, we present an encoding of the diagnosis problem into SAT (Sec. V-C).

A. Preparation

Given the observation α, we need to decide whether all conditions in Theorem 1 hold. Our goal is to minimize the work that is sensitive to changes in α, in particular when α is extended by additional observations. However, (26) and (27) seem to require constructing not only the prefix containing all succinct explanations, but also a large section of the unfolding beyond each of these explanations. Fortunately, this can be avoided thanks to Proposition 4 and the fact that C_1' and C_1 in Theorem 1 do not depend on C being an explanation of α, merely on the fact that C_1' is a fault-free extension of C. So, we only require that $mark(C_1')$ is reachable from $mark(C)$ without executing faults, and replace (27) by:

$$\exists C' \colon mark(C) = mark(C') \wedge C' \subseteq C_1' \tag{32}$$

Hence it suffices to construct only two unfolding prefixes:

- One prefix $\mathcal{P}_\alpha$ containing all succinct explanations of α, used to search for C.
- Prefix $\mathcal{P}^2_N \sqsupseteq \mathcal{P}^1_N$, to check for the existence of a weakly fair configuration starting from a given marking (see Lemma 10), and whether one marking is reachable from another (see Proposition 4). We shall search for C', C_1 in $\mathcal{P}^1_N$ and for C'_1, C_2 in $\mathcal{P}^2_N$.

Observe that restricting the construction of these prefixes to their fault-free parts automatically satisfies (31). Also, notice that $\mathcal{P}_\alpha$ depends only on the observation, whereas $\mathcal{P}^2_N$ only on N. So $\mathcal{P}^2_N$ can be constructed offline, before α is acquired.

B. Constructing the prefixes

We now explain how to compute the prefixes $\mathcal{P}^2_N$ and $\mathcal{P}_\alpha$. There exist well-known algorithms [12] and efficient tools [18], [19] for constructing Petri net unfoldings. Typically, the goal in those constructions is to obtain a marking-complete prefix; to this end, they start with the initial marking of $\mathcal{U}_N$, then discover and add events to the prefix one by one until each branch eventually reaches a so-called *cutoff event* whose causal successors remain unexplored. For our constructions, the iterative structure of these algorithms can be maintained, it suffices to replace the criteria for cutoffs.

1) Constructing $\mathcal{P}_\alpha$: We need to restrict the unfolding construction as follows: (i) exclude fault events to ensure (31); (ii) restrict to explanations of α; and (iii) preserve all succinct explanations and eliminate all verbose ones.

For (i) and (ii), we synchronize N with a net representing $\alpha = \langle S, <, \lambda\rangle$. Let $S_{\min}$ (resp. $S_{\max}$) be the elements without predecessor (resp. successor) in S. We re-translate α into an occurrence net $O_\alpha = \langle P_\alpha, S, F_\alpha, m_\alpha\rangle$, whose events are S and whose causal relation is $<$. The definition of O_α is quite standard, we only remark that $P_\alpha := P_{min} \uplus P_{mid} \uplus P_{max}$ is partitioned in three sets, where P_{max} (resp. P_{min}) is the postset (resp. preset) conditions of $S_{\max}$ (resp. $S_{\min}$).

We then compose $N = \langle P, T, F, m_0\rangle$ and O_α into a net $N_\alpha = \langle P', T_o \cup T_u, G, m'_0\rangle$, where:

- $P' = P \cup P_\alpha$;
- $T_o = \{\, \langle t, s\rangle : t \in T^{obs}, s \in S, \lambda(t) = \lambda(s) \,\}$;
- $T_u = T^{ubs} \setminus \{\phi\}$;
- for $\langle t, s\rangle \in T_o$, ${}^\bullet\langle t, s\rangle = {}^\bullet t \cup {}^\bullet s$ and $\langle t, s\rangle^\bullet = t^\bullet \cup s^\bullet$
- for $t \in T_u$, ${}^\bullet t$ and $t^\bullet$ remain as in N;
- $m'_0 = m_0 \cup m_\alpha$.

Intuitively, N_α adds the places of O_α to N in order to record which parts of α have been seen during an execution. The observable transitions of N and O_α are synchronized to ensure that no run contradict α or add further observable events, and faults are excluded. Consider the unfolding $\mathcal{U}_{N_\alpha}$. Projecting each event labelled with a tuple $\langle t, s\rangle$ to t instead, then each configuration C of $\mathcal{U}_{N_\alpha}$ is also a configuration of $\mathcal{U}_N$; moreover C explains α iff $mark(C)$ contains P_{max}.

It remains to assure (iii). Thanks to Def. 2 it suffices to cut the construction of $\mathcal{U}_{N_\alpha}$ at any event e such that there is another event $e' < e$ with $mark([e']) = mark([e])$. Indeed, this ensures both (9) and (10) as no observable event has occurred after e'. By the pigeon-hole principle on the finitely many reachable marking in N, this cutoff criterion is guaranteed to yield a finite prefix $\mathcal{P}_\alpha$.

2) Constructing $\mathcal{P}^2_N$: We construct $\mathcal{P}^2_N$ in two phases. First, $\mathcal{P}^1_N$ is obtained by the usual unfolding methods for marking-complete prefixes (e.g., [12]). Then, we extend $\mathcal{P}^1_N$ by additional events, using Def. 3 as a cutoff criterion. Observe that deciding whether $e \in E$ is an sp-cutoff entirely depends information contained in $[e]$.

C. Encoding diagnosis into SAT

We propose an encoding of the diagnosis problem into SAT. Given prefixes $\mathcal{P}^1_N, \mathcal{P}^2_N, \mathcal{P}_\alpha$, computed as per Sec. V-B, we construct a formula φ that is satisfiable iff α does not diagnose ϕ. This approach immediately gives a decision procedure via efficient SAT solving. Not surprisingly, one can show (reduction from the reachability problem for unfolding prefixes) that finding the configurations C, C', C_1, C'_1, C_2 discussed in Sec. V-A in these prefixes is NP-hard.

SAT-based decision procedures for unfolding-related problems have been used, e.g. to solve deadlock or reachability problems [20], [21], where satisfying assignments represent configurations with suitable properties. While we re-use this idea, the specificities of diagnosis require to encode multiple configurations and relate them according to Theorem 1.

For an unfolding prefix $\mathcal{P} = \langle\langle B, E, G, \tilde{m}_0\rangle, f\rangle$ of N or N_α, and a *label* l, we define collections of Boolean variables

$$v(l) := \{\mathsf{v}^l_x : x \in B \cup E\}, \qquad m(l) := \{\mathsf{m}^l_p : p \in P\}.$$

Intuitively, all variables in $v(l)$ will encode a configuration identified by l, and those in $m(l)$ a marking referred by l. Write $\mathsf{amo}(\mathsf{v}_1, \ldots, \mathsf{v}_n)$ for 'at most one of $\mathsf{v}_1, \ldots, \mathsf{v}_n$ holds'. For labels l, l', we define the following predicates:

$\star\ \mathsf{config}(l, \mathcal{P}) := \Big(\bigwedge_{e\in E}\bigwedge_{e'\in {}^\bullet{}^\bullet e}(\mathsf{v}^l_e \Rightarrow \mathsf{v}^l_{e'})\Big) \wedge \Big(\bigwedge_{c\in B, \{e_1,\ldots,e_n\}={}^\bullet c} \mathsf{amo}(\mathsf{v}^l_{e_1}, \ldots, \mathsf{v}^l_{e_n})\Big) \wedge \Big(\bigwedge_{c\in B} \mathsf{v}^l_c \Leftrightarrow \big(\bigwedge_{e\in {}^\bullet c} \mathsf{v}^l_e \wedge \bigwedge_{e\in c^\bullet} \neg\mathsf{v}^l_e\big)\Big)$
demands $v(l)$ to represent a configuration of $\mathcal{P}$ (a causally-closed, conflict-free set of events) and its cut.

$\star\ \mathsf{subset}(l, l', \mathcal{P}) := \bigwedge_{e\in E}(\mathsf{v}^l_e \Rightarrow \mathsf{v}^{l'}_e)$ asks that l-labelled events are a subset of l'-labelled events.

$\star\ \mathsf{mark}(l, l', \mathcal{P}) := \bigwedge_{p\in P}\Big(\mathsf{m}^l_p \Leftrightarrow \big(\bigvee_{c\in f^{-1}(p)} \mathsf{v}^{l'}_c\big)\Big)$ asks that $m(l)$ reflects the marking associated with the cut of $v(l)$, assuming that $v(l)$ encodes a configuration of $\mathcal{P}$.

$\star\ \mathsf{enables}(l, l', \mathcal{P}) := \bigwedge_{e\in E}\big(\mathsf{v}^{l'}_e \Leftrightarrow (\bigwedge_{c\in {}^\bullet e} \mathsf{v}^l_c)\big)$ means that, assuming $\mathsf{config}(l, \mathcal{P})$ holds, event variable $\mathsf{v}^{l'}_e$ is true iff e is enabled in the configuration represented by $v(l)$.

$\star\ \mathsf{spoils}(l, l', \mathcal{P}) := \bigwedge_{e\in E}\big(\mathsf{v}^{l'}_e \Rightarrow (\bigvee_{e'\in({}^\bullet e)^\bullet} \mathsf{v}^l_{e'})\big)$ holds iff $v(l)$ has one spoiler for each event true in $v(l')$.

$\star\ \mathsf{explains}(l) := \bigwedge_{p\in P_{max}} \bigvee_{f(c)=p} \mathsf{v}^l_c$, requests that the configuration referred by l is a succinct explanation of α.

Although these predicates are at worst quadratic in the size of $\mathcal{P}$, a linear size version exist for all of them. Also, not all of them are directly given in CNF, but a linear translation is always possible; the same holds for $\mathsf{amo}(\cdot)$, cf. [21].

We can now turn to the encoding of Theorem 1, where (27) is replaced by (32), as discussed in Sec. V-A. Fix labels

$\mathtt{m}, \mathtt{m}', \mathtt{C}, \mathtt{C}', \mathtt{C}'_1, \mathtt{C}_1, \mathtt{C}_2, \mathtt{D}$. Each of these labels identifies (the collection of Boolean variables representing) a configuration or a marking, except for $\mathtt{D}$, which represents a set of events. For instance $\mathtt{C}$ represents the configuration C, note the different typography. Our formula φ is the conjunction of the following constraints:

1) $\mathsf{config}(\mathtt{C}, \mathcal{P}_\alpha) \wedge \mathsf{mark}(\mathtt{m}, \mathtt{C}, \mathcal{P}_\alpha) \wedge \mathsf{explains}(\mathtt{C})$
2) $\mathsf{config}(\mathtt{C}', \mathcal{P}^1_N) \wedge \mathsf{mark}(\mathtt{m}, \mathtt{C}', \mathcal{P}^1_N)$
3) $\mathsf{config}(\mathtt{C}'_1, \mathcal{P}^2_N) \wedge \mathsf{mark}(\mathtt{m}', \mathtt{C}'_1, \mathcal{P}^1_N)$
4) $\mathsf{config}(\mathtt{C}_1, \mathcal{P}^1_N) \wedge \mathsf{mark}(\mathtt{m}', \mathtt{C}_1, \mathcal{P}^2_N)$
5) $\mathsf{config}(\mathtt{C}_2, \mathcal{P}^2_N) \wedge \mathsf{mark}(\mathtt{m}', \mathtt{C}_2, \mathcal{P}^2_N)$
6) $\mathsf{subset}(\mathtt{C}', \mathtt{C}'_1, \mathcal{P}^1_N) \wedge \mathsf{subset}(\mathtt{C}_1, \mathtt{C}_2, \mathcal{P}^2_N)$
7) $\mathsf{enables}(\mathtt{C}_1, \mathtt{D}, \mathcal{P}^2_N) \wedge \mathsf{spoils}(\mathtt{C}_2, \mathtt{D}, \mathcal{P}^2_N)$

Only 1) actually depends on α, whereas remaining constraints can be built before α is known. 7) corresponds to (30), the others to (28), (29), and (32). Conditions (26) and (31) of Theorem 1 are guaranteed by construction.

VI. Outlook

We presented an unfolding-based method solving the problem of *weak diagnosis* for partially observable safe Petri nets.

Weak diagnosis exploits indirect dependencies, captured by the *reveals* relations, in order to explore the system's executions that explain a given observation pattern, and determine whether an unobservable fault is inevitable. We stress that despite its name, 'weak' diagnosis is actually *stronger* than usual diagnosis as in [1]. Whereas in [1] an observation can only be used to detect faults having occurred *in the past*, weak diagnosis captures also faults that are concurrent or in the future of the observation, under weak fairness. The requirement of [1] that no unobservable cycle is present in the system is also dropped, thanks to a characterization of the succinct explanations, that bound the unfolding prefix needed to perform diagnosis. The results here contain those of [2] and strengthen the existing approaches to the more powerful capability of *weak diagnosis*. We have shown how diagnosis can be performed using an algorithmic construction, and given an encoding into SAT.

We intend to produce an implementation of our approach. While many of its ingredients are available by minor modifications of existing tools and verification infrastructure, e.g. [18], a practical obstacle to overcome could be the sizes of the prefixes required in order to perform diagnosis. As for prefix $\mathcal{P}^2_N$, note that it contains all system behaviours, but can be constructed offline once and for all. The size of $\mathcal{P}^2_N$ can be exponential in the size of the net (and polynomial in the size of the reachability graph); however, it is known that unfoldings tend to be much smaller than this for systems that exhibit a high degree of concurrency. In general, the weak-diagnosis problem for Petri nets is PSPACE-complete (hardness follows by reduction from the reachability problem, membership by the fact that a fault-free weakly fair run matching the observation pattern can be nondeterministically simulated in linear space).

The prefix containing all the succinct explanations must be created online, but contains only the behaviours compatible with the observation. Notice that it can alternatively be obtained by producing a marking-complete unfolding of the net N_α from Sec. V-B, which should result in a reduction of its size. In this paper, we omitted this possibility for the sake of a simpler presentation. Also, representing both prefixes as merged processes [20] should result in a dramatic reduction of their size.

Future work also includes verification of *weak diagnosability* [22], [15], [16], [9], [10] based on the results here. Further, local projections of observations, as exploited in [2], are interesting, especially in the context of *distributed* diagnosis.

Acknowledgment: We thank Javier Esparza and the anonymous reviewers for their helpful suggestions. This work has been supported by project ImpRo ANR-2010-BLAN-0317.

References

[1] A. Benveniste, E. Fabre, S. Haar, and C. Jard, "Diagnosis of asynchronous discrete-event systems: a net unfolding approach," *IEEE Trans. Aut. Cont.*, vol. 48, no. 5, pp. 714–727, 2003.
[2] J. Esparza and C. Kern, "Reactive and proactive diagnosis of distributed systems using net unfoldings," in *Proc. ACSD*, 2012, pp. 154–163.
[3] S. Balaguer, T. Chatain, and S. Haar, "Building tight occurrence nets from reveals relations," in *Proc. ACSD*. IEEE, 2011, pp. 44–53.
[4] C. Cassandras and S. Lafortune, *Introduction to discrete event systems*. Springer, 2008.
[5] E. Fabre and A. Benveniste, "Partial order techniques for distributed discrete event systems: Why you cannot avoid using them," *Discrete Event Dynamic Systems*, vol. 17, no. 3, pp. 355–403, 2007.
[6] E. Fabre, A. Benveniste, S. Haar, and C. Jard, "Distributed monitoring of concurrent and asynchronous systems," *Discrete Event Dynamic Systems*, vol. 15, no. 1, pp. 33–84, 2005.
[7] P. Baldan, S. Haar, and B. König, "Distributed unfolding of Petri nets," in *Proc. FoSSaCS*, ser. LNCS 3921, Mar. 2006, pp. 126–141.
[8] S. Haar and É. Fabre, "Diagnosis with Petri net unfoldings," in *Control of Discrete-Event Systems*, ser. LNCIS, C. Seatzu, M. Silva, and J. H. van Schuppen, Eds. Springer, 2013, vol. 433, pp. 301–318.
[9] S. Haar, "Types of asynchronous diagnosability and the reveals-relation in occurrence nets," *IEEE Transactions on Automatic Control*, vol. 55, no. 10, pp. 2310–2320, 2010.
[10] A. Agarwal, A. Madalinski, and S. Haar, "Effective verification of weak diagnosability," in *Proc. 8th SAFEPROCESS*. IFAC, Aug. 2012.
[11] S. Haar, "What topology tells us about diagnosability in partial order semantics," *JDEDS*, vol. 22, no. 4, pp. 383–402, 2012.
[12] J. Esparza, S. Römer, and W. Vogler, "An improvement of McMillan's unfolding algorithm," *Formal Methods in System Design*, vol. 20, no. 3, pp. 285–310, 2002.
[13] W. Vogler, "Fairness and partial order semantics," *Inf. Process. Lett.*, vol. 55, no. 1, pp. 33–39, 1995.
[14] N. Francez, *Fairness*. Springer, 1986.
[15] S. Haar, "Unfold and cover: Qualitative diagnosability for Petri nets," in *Proc. CDC*. IEEE, 2007, pp. 1886–1891.
[16] ——, "Qualitative diagnosability of labeled Petri nets revisited," in *Proc. CDC*. IEEE, 2009, pp. 1248–1253.
[17] S. Haar, C. Kern, and S. Schwoon, "Computing the reveals relation in occurrence nets," in *Proc. GandALF*, ser. ENTCS 54, 2011, pp. 31–44.
[18] S. Schwoon, "MOLE," www.lsv.ens-cachan.fr/~schwoon/tools/mole/.
[19] V. Khomenko, "PUNF," http://homepages.cs.ncl.ac.uk/victor.khomenko/tools/punf/.
[20] V. Khomenko, A. Kondratyev, M. Koutny, and W. Vogler, "Merged processes – a new condensed representation of Petri net behaviour," *Acta Informatica*, vol. 43, no. 5, pp. 307–330, 2006.
[21] C. Rodríguez and S. Schwoon, "Verification of Petri nets with read arcs," in *Proc. CONCUR*, ser. LNCS 7454, 2012, pp. 471–485.
[22] S. Haar, A. Benveniste, E. Fabre, and C. Jard, "Partial order diagnosability of discrete event systems using Petri net unfoldings," in *Proc. CDC*. IEEE, 2003, pp. 3748–3753.