

Basis Coverability Graph for Partially Observable Petri Nets with Application to Diagnosability Analysis

Engel Lefauchaux, Alessandro Giua, Carla Seatzu

► To cite this version:

Engel Lefauchaux, Alessandro Giua, Carla Seatzu. Basis Coverability Graph for Partially Observable Petri Nets with Application to Diagnosability Analysis. Petri Nets 2018 - International Conference on Applications and Theory of Petri Nets and Concurrency, Jun 2018, Bratislava, Slovakia. pp.164-183, 10.1007/978-3-319-91268-4_9. hal-01882129

HAL Id: hal-01882129

<https://hal.inria.fr/hal-01882129>

Submitted on 26 Sep 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Basis Coverability Graph for Partially Observable Petri Nets with Application to Diagnosability Analysis

Engel Lefauchaux^{a,b}, Alessandro Giua^{c,d}, Carla Seatzu^c

a. Univ. Rennes, INRIA, Campus Universitaire de Beaulieu, Rennes, France

b. LSV, ENS Paris-Saclay, CNRS, Cachan, France

c. Dept. of Electrical and Electronic Engineering, University of Cagliari, Cagliari, Italy

d. Aix Marseille Univ, Université de Toulon, CNRS, ENSAM, LSIS, Marseille, France

Abstract. Petri nets have been proposed as a fundamental model for discrete-event systems in a wide variety of applications and have been an asset to reduce the computational complexity involved in solving a series of problems, such as control, state estimation, fault diagnosis, etc. Many of those problems require an analysis of the reachability graph of the Petri net. The basis reachability graph is a condensed version of the reachability graph that was introduced to efficiently solve problems linked to partial observation. It was in particular used for diagnosis which consists in deciding whether some fault events occurred or not in the system, given partial observations on the run of the system. However this method is, with very specific exceptions, limited to bounded Petri nets. In this paper, we introduce the notion of basis coverability graph to remove this requirement. We then establish the relationship between the coverability graph and the basis coverability graph. Finally, we focus on the diagnosability and stochastic diagnosability problems: we show how the basis coverability graph can be used to get efficient algorithms when such problems are decidable.

Introduction

The *marking reachability problem* is a fundamental problem of Petri nets (PNs) which can be stated as follows: *Given a net system $\langle N, M_0 \rangle$ and a marking M , determine if M belongs to the reachability set $R(N, M_0)$.* It plays an important role since many other properties of interest can be solved by reduction to this problem. The marking reachability problem has been shown to be decidable in [19] and was shown to be EXPSPACE-hard in [23].

In the case of *bounded* PNs, i.e., net systems whose reachability set is finite, a straightforward approach to solve this problem consists in constructing the *reachability graph*, which provides an explicit representation of the net behavior, i.e., its reachability set and the corresponding firing sequences of transitions. However, albeit finite, the reachability graph may have a very large number of nodes due to the so called *state space explosion* that originates from the combinatorial nature of discrete event systems. For this reason, practically efficient

approaches, which do not require to generate the full state space, have been explored. We mention, among others, partial order reduction techniques, such as the general approaches based on stubborn sets [28] and persistent sets [13] or the Petri net approaches based on unfolding [21] and maximal permissive steps [4].

In the case of *unbounded* PNs, whose reachability set is infinite, the authors of [16] have shown that a finite *coverability graph* may be constructed which provides a semi-decision procedure (necessary conditions) for the marking reachability problem. It provides an over-approximation of both the reachability set and the set of firing sequences. As was the case for the reachability graph, this approach is not efficient and improvements to the basic algorithm have later been proposed [22].

Recently some of us have proposed a quite general approach that exploits the notion of *basis marking* to practically reduce the computational complexity of solving the reachability problem for bounded nets. This method has originally been introduced to solve problems of state estimation under partial observation [12] but has later been extended to address fault diagnosis [6], state-based opacity [27] and general reachability problems [17].

The approach in [17] considers a partition of the set of transitions $T = T_e \cup T_i$: T_e is called set of *explicit transitions* and T_i is called set of *implicit transitions*. The main requirement is that the subnet containing only implicit transitions be acyclic. The firing of implicit transitions is abstracted and only the firing of explicit transitions need to be enumerated. The advantage of this technique is that only a subset of the reachability space — i.e., the set of the so-called basis markings — is enumerated. All other markings are reachable from a basis marking by firing only implicit transitions and can be characterized by the integer solutions of a system of linear equations. In a certain sense, this hybrid approach combines a behavioral analysis (limited to the firing of transitions in T_e) with a structural analysis (which describes the firing of transitions in T_e).

The objective of this paper is twofold. First, we show that the approach of [17] can be generalized to unbounded nets and we define a *basis coverability graph* where the firing of implicit transitions is abstracted, thus reducing the number of nodes of the standard coverability graph. In addition, we show how this approach can be applied to study the *diagnosability* of Petri nets in the logic framework of [24] and in the stochastic framework of [2]. In this case, we consider as implicit the set of unobservable transitions. However, since the firing of unobservable faulty transitions need to be recorded, we further extend the approach of [17] by considering that there may exist a subset of implicit transitions (called *relevant transitions*) which, albeit abstracted, need to be handled with special care. The diagnosis of both bounded and unbounded nets is considered.

The paper is structured as follows. In Section 1, we recall some usual definitions for Petri Nets and their coverability graph. In Section 2, we introduce the notion of basis coverability graph and establish some of its properties. In Section 3 we give the definitions of stochastic Petri nets and of logical/stochastic diagnosability. In Section 4 we show how to use the basis coverability graph to analyse the stochastic diagnosability of bounded Petri nets. Finally in Section 5 we study

unbounded Petri nets: we prove the undecidability of the stochastic diagnosability analysis and how to use the basis reachability graph for the logical diagnosability analysis.

Due to space constraints, one of the results presented in the paper is provided without proof. For the convenience of the reviewers, the proof is contained in an appendix which will be removed if the paper is accepted for publication. In the final version a link to a full version on HAL will be provided instead.

1 Background on Petri nets and Coverability Graph

1.1 Petri Nets

In this section the formalism used in the paper is recalled. For more details on Petri nets the reader is referred to [20].

Definition 1. A Petri net (PN) is a structure $N = (P, T, Pre, Post)$, where P is a set of m places; T is a set of n transitions; $Pre : P \times T \rightarrow \mathbb{N}$ and $Post : P \times T \rightarrow \mathbb{N}$ are the pre- and post- incidence functions that specify the arcs. We also define $C = Post - Pre$ as the incidence matrix of the net.

A marking is a vector $M : P \rightarrow \mathbb{N}$ that assigns to each place of a PN a nonnegative integer number of tokens. A net system (NS) $\langle N, M_0 \rangle$ is a PN N with an initial marking M_0 . A transition t is enabled at M iff $M \geq Pre(\cdot, t)$ and may fire yielding the marking $M' = M + C(\cdot, t)$. One writes $M [\sigma]$ to denote that the sequence of transitions $\sigma = t_{j_1} \cdots t_{j_k}$ is enabled at M , and $M [\sigma] M'$ to denote that the firing of σ yields M' . One writes $t \in \sigma$ to denote that a transition t is contained in σ . The length of the sequence σ (denoted $|\sigma|$) is the number of transitions in the sequence, here k .

The set of all sequences that are enabled at the initial marking M_0 is denoted $L(N, M_0)$, i.e., $L(N, M_0) = \{\sigma \in T^* \mid M_0[\sigma]\}$. Given $k \geq 0$, the set of all sequences of length k is written T^k , the set of all infinite sequences is written T^ω . A marking M is *reachable* in $\langle N, M_0 \rangle$ iff there exists a firing sequence σ such that $M_0 [\sigma] M$. The set of all markings reachable from M_0 defines the *reachability set* of $\langle N, M_0 \rangle$ and is denoted $R(N, M_0)$.

Let $\pi : T^* \rightarrow \mathbb{N}^n$ be the function that associates with the sequence $\sigma \in T^*$ a vector $y \in \mathbb{N}^n$, called the *firing vector* of σ . In particular, $y = \pi(\sigma)$ is such that $y(t) = k$ iff the transition t is contained k times in σ .

A PN having no directed circuits is called *acyclic*. Given $k \in \mathbb{N}$, a place p of an NS $\langle N, M_0 \rangle$ is *k-bounded* if for all $M \in R(N, M_0)$, $M(p) \leq k$. It is bounded if there exists $k \in \mathbb{N}$ such that p is *k-bounded*. An NS is bounded (resp. *k-bounded*) iff all of its places are bounded (resp. *k-bounded*).

A sequence is *repetitive* iff it can be repeated indefinitely (i.e. σ is repetitive in the marking M iff $M[\sigma]M'$ with $M' \geq M$). There are two kinds of repetitive sequences: a repetitive sequence is *stationary* if it does not modify the marking (i.e. $M[\sigma]M$), it is *increasing* otherwise. Remark that an NS containing an increasing sequence can not be bounded.




Fig. 1. A net system. Circles are places and rectangles are transitions. In the initial marking, p_1 has two tokens represented by the two black dots.

Example 1. Consider the NS of Figure 1, the sequence t_1 , is increasing in the initial marking $M_0 = [2, 0, 0, 0, 0]$. Firing t_1 k times in M_0 leads to the marking $M_1 = [2, k, 0, 0, 0]$. Therefore the place p_2 is not bounded. However, every other place is 2-bounded.

Definition 2. Given a net $N = (P, T, Pre, Post)$, and a subset $T' \subseteq T$ of its transitions, let us define the T' -induced subnet of N as the new net $N' = (P, T', Pre', Post')$ where $Pre', Post'$ are the restrictions of $Pre, Post$ to T' . The net N' can be thought as obtained from N removing all transitions in $T \setminus T'$. Let us also write $N' \prec_{T'} N$.

1.2 Coverability Graph

For a bounded NS $\langle N, M_0 \rangle$, one can enumerate the elements of the reachability set $R(N, M_0)$ and establish the transition function between the markings. The resulting graph is called *Reachability Graph*. If the NS is not bounded, this construction does not terminate. Instead, an usual method is to build the *Coverability Graph* which is a finite over-approximation of the reachability set and of the net language [16]. We will define in this section the coverability graph of an NS which if the NS is bounded is equal to the reachability graph of this NS.

An ω -marking is a vector from the set of places to $\mathbb{N} \cup \{\omega\}$, where ω should be thought of as "arbitrarily large": for all $k \in \mathbb{N}$, we have $k < \omega$ and $\omega \pm k = \omega$. An ω -marking M is (resp. strictly) covered by an ω -marking M' , written $M \leq M'$ (resp. $M \prec M'$) iff for every place p of the net, $M(p) \leq M'(p)$ (resp. and there exists at least one place p such that $M(p) < M'(p)$).

Definition 3. Given an NS $\langle N, M_0 \rangle$, the associated coverability graph $CG_{\langle N, M_0 \rangle} = (\mathcal{M}, M_0, \Delta)$ is defined in the following manner.

We first define inductively a temporary set \mathcal{M}_t of pairs of ω -markings and set of ω -markings and the temporary transition function Δ_t by:

- $(M_0, \{M_0\}) \in \mathcal{M}_t$ and
- $(M', B') \in \mathcal{M}_t$ iff there exists $(M, B) \in \mathcal{M}_t$ and $t \in T$ such that
 - either $M[t]M', B' = B \cup \{M'\}$ and for all $M'' \in B, M' \not\prec M''$;

- or, for M^t such that $M[t]M^t$, there exists $M'' \in B$ such that $M^t \succeq M''$. For every such M'' , let p_1, \dots, p_k be the set of places such that for all j , $M^t(p_j) \succeq M''(p_j)$, then $\forall j, M'(p_j) = \omega$. For every place p such that $M'(p) \neq \omega$, $M'(p) = M^t(p)$. Moreover $B' = B \cup \{M'\}$.

In both cases, $((M, B), t, (M', B')) \in \Delta_t$.

We then define $\mathcal{M} = \{M \mid \exists B, (M, B) \in \mathcal{M}_t\}$ and given M and M' in \mathcal{M} , $(M, t, M') \in \Delta$ iff there exists B, B' such that $((M, B), t, (M', B')) \in \Delta_t$.

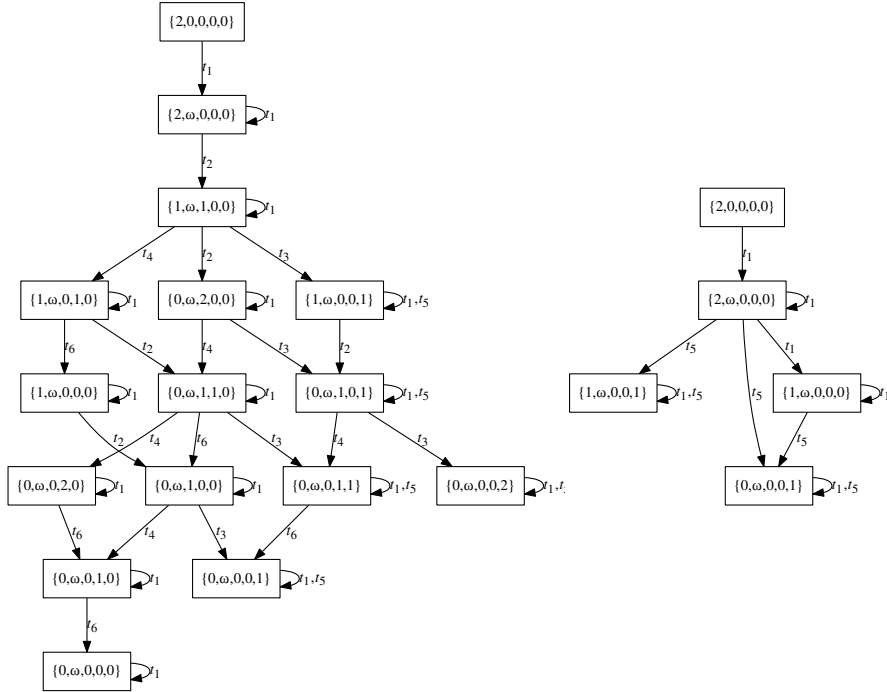


Fig. 2. Left: the coverability graph of the NS in Figure 1. Right: The BCG of the NS in Figure 1 where $T_i = \{t_2, t_3, t_4, t_6\}$ and $T_s = \{t_5\}$. The firing vectors are omitted on the edges.

The temporary graph built here is equivalent to the coverability tree of [7]. They proved in [16] that the coverability tree (and thus our temporary graph) terminates in a finite number of steps.

Example 2. The coverability graph of the NS in Figure 1 is shown in Figure 2. The firing of t_1 at the initial marking adds a token to the second place, reaching a marking strictly greater than the initial marking in this place and equal everywhere else. Correspondingly in the coverability graph an ω appears in the second component of the marking to show that there is a repetitive sequence enabled by the system which increases the number of tokens in the second place.

A marking M is ω -covered by an ω -marking M_ω , denoted $M \leq_\omega M_\omega$ if for every place p such that $M_\omega(p) \neq \omega$, $M_\omega(p) = M(p)$. Using this definition and the coverability graph, we define the coverability set of an NS which is an over-approximation of the reachability set.

Definition 4. Given an NS $\langle N, M_0 \rangle$, let \mathcal{M} be the set of ω -markings of its coverability graph, the coverability set of $\langle N, M_0 \rangle$ is

$$CS(N, M_0) = \{M \in \mathbb{N}^m \mid \exists M_\omega \in \mathcal{M}, M \leq_\omega M_\omega\}$$

Example 3. The coverability set of the NS in Figure 1 is equal to its reachability set. This is not the case however for the NS in Figure 3 where the reachability set is $\{(k, 2r) \mid k, r \in \mathbb{N}\}$ while the coverability set is $\{(k, r) \mid k, r \in \mathbb{N}\}$. We however clearly see that the coverability set subsumes the reachability set.




Fig. 3. A Petri net where the coverability set strictly subsumes the reachability set. Transition t_2 is unobservable.

We will use the rest of this section to recall a few known applications of the coverability graph and the coverability set. All those results can be found in [7]. First, as claimed earlier, the coverability set subsumes the reachability set.

Proposition 1. Let $\langle N, M_0 \rangle$ be an NS, $R(N, M_0) \subseteq CS(N, M_0)$.

The coverability graph can be used to determine if an NS is bounded.

Proposition 2. Given an NS $\langle N, M_0 \rangle$,

- a place p is k -bounded \Leftrightarrow for each marking M of $CG_{\langle N, M_0 \rangle}$, $M[p] \leq k$.
- the marked net is bounded \Leftrightarrow no node of $CG_{\langle N, M_0 \rangle}$ contains the symbol ω

Repetitiveness can be partially checked on the coverability graph.

Proposition 3. Given an NS $\langle N, M_0 \rangle$, a marking M and a non-empty sequence σ' of transitions enabled by M ,

- σ is repetitive \Rightarrow there exists a directed cycle in the coverability graph whose arcs form σ starting in an ω marking M_ω such that $M_\omega \geq_\omega M$.
- σ is stationary \Leftrightarrow there exists a directed cycle starting in M in the graph that does not pass through markings containing ω and whose arcs form σ .

2 Basis coverability graph

While the reachability/coverability graph has many applications, one of its downside is its size. For bounded NS, the authors of [6, 9] introduced the notion of basis reachability graph which keeps most of the information relevant for partially observed systems of the reachability graph while decreasing, in some cases exponentially, the size of the graph. Their goal at the time was to study diagnosis. They then generalised this approach to study reachability (regardless of labeling on transitions) in [17]. The idea of the basis reachability graph is to select a set of transitions called "implicit" in [17] (and unobservable in [6]) that will be abstracted and to only represent the "explicit" transitions that can be fired (possibly after some implicit transition) in a given marking. In this section, we will describe how to apply this idea to unbounded NS and how to build instead a *Basis Coverability Graph* (BCG). When the NS is bounded, the BCG is equal to the basis reachability graph.

Given a set of transitions T of a PN, we denote $T_i \subseteq T$ and $T_e = T \setminus T_i$ the sets of *implicit* and *explicit* transitions respectively. Let C_i (C_e) be the restriction of the incidence matrix to T_i (T_e) and n_i and n_e , respectively, be the cardinality of the above sets of transitions. Given a sequence $\sigma \in T^*$, $P_i(\sigma)$, resp., $P_e(\sigma)$, denotes the projection of σ over T_i , resp., T_e .

We will sometimes need the following assumptions.

A1: The T_i -induced subnet is acyclic.

A2: Every sequence containing only implicit transitions is of finite length.

Remark that for bounded NS, the first assumption, which is an usual requirement for problems such as diagnosis of discrete event systems, implies the second one.

When the partition between implicit and explicit transitions is not given, one can always choose a partition respecting the two assumptions above (for example $T_e = T$). The authors of [17] discuss how to choose an appropriate partition for the basis reachability graph and how this choice affects the cardinality of the set of markings of the graph.

Definition 5. Given a marking M and an explicit transition $t \in T_e$, let

$$\Sigma(M, t) = \{\sigma \in T_i^* \mid M[\sigma]M', M' \geq \text{Pre}(\cdot, t)\}$$

be the set of *explanations* of t at M , and let

$$Y(M, t) = \pi(\Sigma(M, t))$$

be the *e-vectors* (or *explanation vectors*), i.e., firing vectors associated with the explanations.

Thus $\Sigma(M, t)$ is the set of implicit sequences whose firing at M enables t . Among the above sequences we will select those whose firing vector is minimal and those who are minimal while containing a transition among a chosen set

$T_s \subseteq T_i$ which will be called the set of *relevant* transitions. This second category is used to solve problems where it may be necessary to keep track of a subset of implicit transitions. In particular it will be used in the sections about diagnosis later in this paper. The firing vector of these sequences are called (T_s -) *minimal e-vectors*.

Definition 6. Given a marking M , a transition $t \in T_e$ and a set of relevant transitions $T_s \subseteq T_i$, let

$$\Sigma_{\min}(M, t) = \{\sigma \in \Sigma(M, t) \mid \nexists \sigma' \in \Sigma(M, t) : \pi(\sigma') \preceq \pi(\sigma)\}$$

be the set of *minimal explanations* of t at M ,

$$\Sigma_{\min}^{T_s}(M, t) = \{\sigma \in \Sigma(M, t) \mid \sigma \cap T_s \neq \emptyset \wedge \nexists \sigma' \in \Sigma(M, t) : \sigma \cap T_s = \sigma' \cap T_s \wedge \pi(\sigma') \preceq \pi(\sigma)\}$$

the set of T_s -*minimal explanations* of t at M .

Remark that for two sets of relevant transitions $T_s \subseteq T_i$ and $T'_s \subseteq T_i$, if $T_s \subseteq T'_s$, for every marking M and explicit transition $t \in T_e$, $\Sigma_{\min}^{T_s}(M, t) \subseteq \Sigma_{\min}^{T'_s}(M, t)$. We will now build the BCG with a construction similar to the one of the coverability graph using minimal explanations instead of transitions.

Definition 7. Given an NS $\langle N, M_0 \rangle$ and a set of relevant transition $T_s \subseteq T_i$, the associated basis coverability graph (BCG) with relevant transitions T_s $BCG_{\langle N, M_0 \rangle}^{T_s} = (\mathcal{M}, M_0, \Delta)$ is defined in the following manner.

We first define inductively a temporary set \mathcal{M}_t of pairs of ω -markings and set of ω -markings and the temporary transition function Δ_t by:

- $(M_0, \{M_0\}) \in \mathcal{M}_t$ and
- $(M', B') \in \mathcal{M}_t$ iff there exists $(M, B) \in \mathcal{M}_t$, $t \in T_e$ and $\sigma = t_1 \dots t_n \in \Sigma_{\min}(M, t) \cup \Sigma_{\min}^{T_s}(M, t)$ with $M[t_1]M_1[t_2] \dots [t_n]M_n[t]M_{n+1}$ such that
 - either $M_{n+1} = M'$, $B' = B \cup \{M_1, \dots, M_{n+1}\}$ and for all $M'' \in B$, $i \in \{1, \dots, n+1\}$, $M_i \not\geq M''$;
 - or there exists $i \geq 1$ and $M'' \in B \cup \{M_j \mid j < i\}$ such that $M_i \geq M''$. First, we pick the minimal such i . Then, for every $M'' \in B \cup \{M_j \mid j < i\}$ such that $M_i > M''$, let p_1, \dots, p_k be the set of places such that for all $1 \leq j \leq k$, $M_i(p_j) \geq M''(p_j)$. We transform the markings M_r , $r \geq i$ into ω -markings where for all $1 \leq j \leq k$, $M_r(p_j) = \omega$, the other places being unchanged. We then repeat this process on the obtained sequence of markings until there is no i verifying the property. Let M'_1, \dots, M'_{n+1} be the result of the repeated process. Then $B' = B \cup \{M'_1, \dots, M'_n\}$ and $M' = M'_n$. Remark that the path $M'[t_1]M'_1[t_2] \dots [t_n]M'_n$ is a path of the coverability graph.

In both cases $((M, B), (\pi(\sigma), t), (M', B')) \in \Delta_t$.

We then define $\mathcal{M} = \{M \mid \exists B, (M, B) \in \mathcal{M}_t\}$ and given M and M' in \mathcal{M} , $(M, (\pi(\sigma), t), M') \in \Delta$ iff there exists B, B' such that $((M, B), (\pi(\sigma), t), (M', B')) \in \Delta_t$.

The markings of the BCG are called *basis markings*.

We denote the projection of a sequence $\sigma = (\sigma_1, t_1) \dots (\sigma_k, t_k) \dots$ of the BCG on its second component by $P_t(\sigma) = t_1 \dots t_k \dots$.

Example 4. We represent the BCG of the NS in Figure 1 (for $T_i = \{t_2, t_3, t_4, t_6\}$ and $T_s = \{t_6\}$. For readability the firing vectors on the edges are omitted in the figure) in Figure 2. This BCG has 11 less states than the coverability graph.

Choosing $T_s = \{t_6\}$ adds the states $\{1, \omega, 0, 0, 0\}$ and $\{0, \omega, 0, 0, 1\}$ and the edges affecting those states. The edge from $\{2, \omega, 0, 0, 0\}$ to $\{0, \omega, 0, 0, 1\}$ corresponds to the $\{t_6\}$ -minimal explanation $t_2 t_2 t_3 t_4 t_6$ of t_5 which is not a minimal explanation.

As hinted to in the example, the BCG is smaller than the coverability graph. This is formally proved in the following.

Proposition 4. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i , for any set of relevant transitions $T_s \subseteq T_i$ it holds that every basis marking M of $BCG_{\langle N, M_0 \rangle}^{T_s}$ is a marking of $CG_{\langle N, M_0 \rangle}$.*

Proof. As any marking of the BCG is reachable from M_0 , we will show the result by induction on the length of a path reaching this marking. Let M be a marking of $BCG_{\langle N, M_0 \rangle}^{T_s}$ and σ a sequence such that $M_0[\sigma]M$.

If $|\sigma| = 0$, $M = M_0$ which is a marking of $CG_{\langle N, M_0 \rangle}$.

Given $n \in \mathbb{N}$, suppose that the property is true for every marking reached by a path of length at most n . If $|\sigma| = n + 1$, $\sigma = \sigma_1(e, t)$, let M_1 such that $M_0[\sigma_1]M_1$, then by hypothesis M_1 belongs to $CG_{\langle N, M_0 \rangle}$. Moreover, as there is a transition $M_1[(e, t)]M$ in the BCG, there exists a minimal explanation $\sigma' = t_1, \dots, t_n \in \Sigma_{\min}(M_1, t) \cup \Sigma_{\min}^{T_s}(M_1, t)$ such that $\pi(\sigma') = e$ and one of the two conditions for a BCG transition between M_1 and M is validated. By definition of those two conditions, there is a path $M_1[t_1]M_2 \dots [t_n]M_n[t]M$ in $CG_{\langle N, M_0 \rangle}$ as remarked in the definition of the BCG, thus M is a marking of $CG_{\langle N, M_0 \rangle}$. \square

We will now give some results showing that the BCG can effectively be used in many cases instead of the coverability graph. As a first step, we will show that the BCG can be used to define a set of markings that are an over-approximation of the reachability set. We denote by $R_i(N, M)$ the set of markings reachable from M using only implicit transitions in the Petri net N . Given an ω -marking M_ω and a marking M , $M_\omega =_\omega M$ iff for every place p such that $M_\omega(p) \neq \omega$, $M_\omega(p) = M(p)$.

Definition 8. *Given an NS $\langle N, M_0 \rangle$ with m places and a set of implicit transitions T_i , let $T_s \subseteq T_i$ be a set of relevant transitions and let V be the set of basis markings of $BCG_{\langle N, M_0 \rangle}^{T_s}$. The basis coverability set of $\langle N, M_0 \rangle$ with relevant transitions T_s is*

$$BCS^{T_s}(N, M_0) = \{M \in \mathbb{N}^m \mid \exists M_\omega \in V, \exists M_\omega^u \in R_i(N, M_\omega), M_\omega^u \geq_\omega M\}$$

This set can be easily computed for NS verifying (A1). For every possible choice of T_s , the basis coverability set is an over-approximation of the reachability set.

Proposition 5. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i verifying Assumption (A1) and a set of relevant transitions $T_s \subseteq T_i$, it holds $R(N, M_0) \subseteq BCS^{T_s}(N, M_0)$.*

Proof. Let σ be a sequence such that $M_0[\sigma]M$ in the NS. We will proceed by induction on the length of σ .

If $|\sigma| = 0$, $M = M_0$ which is a marking of the BCG.

Given $n \in \mathbb{N}$, supposing that the property is true for every marking reached by a path of length at most n . For $|\sigma| = n + 1$, $\sigma = \sigma_1 t$. Let $M_0[\sigma_1]M_1$. By the induction hypothesis there exists a basis marking M_ω^b and an ω -marking M_ω^u such that $M_\omega^u \in R_i(N, M_\omega^b)$ and $M_\omega^u \geq_\omega M_1$.

- if t is implicit, as $M_\omega^u \geq_\omega M_1$, t is enabled by M_ω^u and the marking reached by firing t in M_ω^u , let's call it $M_\omega^{u,2}$, verifies $M_\omega^{u,2} \geq_\omega M$ and $M_\omega^{u,2} \in R_i(N, M_\omega^b)$.
- if t is explicit, let σ_t such that $M_\omega^b[\sigma_t]M_\omega^u$. Since σ_t is an explanation of t , there thus exist a minimal explanation σ_{min} such that $\pi(\sigma_{min}) \leq \pi(\sigma_t)$ and a sequence $\sigma_e \in T_i^*$ such that $\pi(\sigma_{min}) + \pi(\sigma_e) = \pi(\sigma_t)$. Let $M_s \leq_\omega M_\omega^b$ such that $M_s[\sigma_t]M_1$ and M_f the marking such that $M_s[\sigma_{min}t]M_f$. Using Theorem 3.8 of [6] which requires (A1) (in fact this is the result used everytime (A1) is required in the following), $M_f[\sigma_e]M$. By construction of the BCG, there exists a basis marking M_2^b reachable with transition $(\pi(\sigma_{min}), t)$ from M_ω^b such that $M_2^b \geq_\omega M_f$. Moreover, as $M_f[\sigma_e]M$, triggering σ_e in M_2^b leads to a marking M_ω^2 such that $M_\omega^2 \geq_\omega M$. \square

The following result characterizes a monotonicity property of the basis coverability set with respect to the corresponding set of relevant transitions.

Proposition 6. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i . For any two sets of relevant transitions T_s and T'_s such that $T_s \subseteq T'_s \subseteq T_i$, $BCS^{T_s}(N, M_0) \subseteq BCS^{T'_s}(N, M_0)$.*

Proof. Let $\langle N, M_0 \rangle$ be an NS and T_s and T'_s be two sets of implicit transitions such that $T_s \subseteq T'_s$. Let $M \in BCS^{T_s}(N, M_0)$ there thus exists a basis marking M_ω of $BCG_{\langle N, M_0 \rangle}^{T_s}$ such that there exists an ω -marking $M_\omega^u \in R_i(N, M_\omega)$, with $M_\omega^u \geq_\omega M$. As $\Sigma_{\min}^{T_s}(M, t) \subseteq \Sigma_{\min}^{T'_s}(M, t)$, by construction of the BCG, every basis marking of $BCG_{\langle N, M_0 \rangle}^{T_s}$ is a basis marking of $BCG_{\langle N, M_0 \rangle}^{T'_s}$. Thus M_ω is a basis marking of $BCG_{\langle N, M_0 \rangle}^{T'_s}$. Therefore $M \in BCS^{T'_s}(N, M_0)$. \square

The inclusion can be strict. Indeed, let us observe the NS of Figure 3 with t_2 implicit. The BCG with $T_s = \emptyset$ has two basis markings $[0, 0]$ and $[\omega, 0]$. The associated basis coverability set is $\{[n, 2m] \mid n, m \in \mathbb{N}\}$, which is equal to the reachability set. However, the BCG with $T_s = \{t_2\}$ has the two previous basis markings plus $[\omega, \omega]$. Therefore its basis coverability set is $\{[n, m] \mid n, m \in \mathbb{N}\}$, which is equal to the coverability set. In fact the basis coverability set is always a better approximation than the coverability set.

Proposition 7. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i and a set of relevant transitions $T_s \subseteq T_i$, it holds $BCS^{T_s}(N, M_0) \subseteq CS(N, M_0)$.*

Proof. Let $M \in BCS(N, M_0)$. By definition, there exists $M' \geq_\omega M$ and M_b state of the BCG with $M' \in R_i(N, M_b)$. By Proposition 4, M_b is a state of $CG_{\langle N, M_0 \rangle}$. Let σ be an implicit sequence such that $M_b[\sigma]M'$. By definition of R_i and of the coverability graph there exists a state M_c of $CG_{\langle N, M_0 \rangle}$ such that $M_b[\sigma]M_c$ in the CG and $M_c \geq M'$. Thus $M_c \geq M$. Therefore $M \in CS(N, M_0)$. \square

We now show how the results relative to the coverability graph recalled in the previous section (namely Propositions 2 and 3) can be transposed on the BCG. As those results hold for every choice of set of relevant transitions $T_s \subseteq T_i$, this set is omitted for the rest of the section.

Proposition 8. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i verifying Assumptions (A1) and (A2),*

- *a place p is k -bounded \Rightarrow for each basis marking M of the BCG $M[p] \leq k$. The reverse implication is false, one would need to analyse the implicit reach as stated in the next item;*
- *p is not k -bounded \Rightarrow there exists a basis marking M_ω and an ω -marking $M_u \in R_i(N, M_\omega)$ with $M_u(p) > p$;*
- *the NS is bounded \Leftrightarrow no basis marking of the BCG contains the symbol ω .*

Proof. – The first item holds as this implication is true for every marking of the coverability graph according to Proposition 2 and Proposition 4 which claims that the markings of the BCG are markings of the coverability graph. Moreover the inverse implication does not hold: observing the NS of Figure 3, the two basis markings of the BCG are $[0, 0]$ and $[\omega, 0]$, however the second place is not bounded by 0, in fact it is not bounded at all. In this respect, the BCG may not explicitly show all the informations that appears in the coverability graph.

- Suppose that p is not k -bounded. There thus exists a marking $M \in R(N, M_0)$ with $M(p) > k$. As $R(N, M_0) \subseteq BCS(N, M_0)$ according to Proposition 5, there exists a basis marking M_ω and an ω -marking $M_u \in R_i(N, M_\omega)$ such that $M_u \geq_\omega M$. Thus $M_u(p) > k$.
- For the third item, the left to right implication is once again due to Proposition 4 and the fact that the equivalence holds when considering the coverability graph as stated in Proposition 2.

For the right to left implication, suppose that no ω appears in the BCG. Then $BCS(N, M_0)$ is finite as in every basis marking, which are also markings reachable in the NS, there is finitely many sequences of implicit transitions enabled thanks to assumption (A2). Therefore, according to Proposition 5 the reachability set $R(N, M_0)$ is finite. This implies that the NS is bounded. \square

Proposition 9. *Given an NS $\langle N, M_0 \rangle$ with set of implicit transitions T_i verifying Assumption (A1), a non-empty sequence σ' of explicit transitions and a marking M*

- *there exists a repetitive sequence σ with $P_e(\sigma) = \sigma'$ enabled by $M \Rightarrow$ there exists $k \in \mathbb{N}$, two basis markings M_ω^1, M_ω^2 and two ω -markings $M_u^i \in R_i(N, M_\omega^i)$, $i \in \{1, 2\}$, such that:*

- $M \leq_\omega M_u^i, i \in \{1, 2\}$;
 - there is a path starting in M_ω^1 and ending in M_ω^2 in the BCG whose arcs, projected on the second component, form σ' ;
 - there is a directed cycle starting in M_ω in the BCG whose arcs, projected on the second component, form $(\sigma')^k$.
- there exists a directed cycle starting in M_ω in the BCG that does not pass through markings containing ω and whose arcs, projected on the second component, form σ' where M_ω is a basis marking such that $M \in R_i(N, M_\omega) \Rightarrow$ there exists a stationary sequence σ with $P_e(\sigma) = \sigma'$ enabled by M .

Proof. – Suppose that σ is repetitive from M . Due to Proposition 5, there exists a basis marking M_ω^0 and an ω -marking M_u^0 with $M_u^0 \geq_\omega M$ and $M_u^0 \in R_i(N, M_\omega^0)$. Let $\sigma = \sigma_1 t_1 \dots \sigma_n t_n \sigma_{n+1}$ where the σ_i 's are sequences of implicit transitions and the t_i 's are explicit transitions. As the NS verifies (A1) and by construction of the BCG, there exists a sequence $\sigma^1 = \sigma_1^1 t_1 \dots \sigma_n^1 t_n$ enabled by M_ω^0 where the σ_i^1 's are minimal explanations of the t_i 's and ending in a basis marking M_ω^1 such that there exists an ω -marking M_u^1 with $M_u^1 \geq_\omega M$ and $M_u^1 \in R_i(N, M_\omega^1)$. This translates in the BCG into a sequence $(\pi(\sigma_1^1), t_1) \dots (\pi(\sigma_n^1), t_n)$ from M_ω^0 to M_ω^1 . This can be repeated, giving a family of sequences $(\sigma^j)_{j \in \mathbb{N}}$, of basis markings $(M_\omega^j)_{j \in \mathbb{N}}$ and of ω -marking $(M_u^j)_{j \in \mathbb{N}}$ such that $M_\omega^{j-1}[\sigma^j]M_\omega^j, M_u^j \geq_\omega M$ and $M_u^j \in R_i(N, M_\omega^j)$. Due to the finite number of basis markings, there exists $k, k', k < k'$, such that $M_\omega^k = M_\omega^{k'}$. There thus exists a directed cycle starting in M_ω^k whose arcs, projected on the second component, form $P_e(\sigma)^{k'-k}$.

– Suppose that there exists a directed cycle starting in the basis marking M_ω in the BCG that does not pass through markings containing ω and whose arcs, projected on the second component, form σ' . Using the Proposition 4, M_ω is a marking of $CG_{\langle N, M_0 \rangle}$. Moreover due to the construction of the BCG there exists σ such that $P_e(\sigma) = \sigma'$ and a directed cycle starting in M_ω in $CG_{\langle N, M_0 \rangle}$ that does not pass through markings containing ω and whose arcs form σ . Due to Proposition 3, this implies that σ is stationary. \square

3 Stochastic Petri Nets and Diagnosability

Probabilities are added to an NS by adding a fire rate to every transition in the following way.

Definition 9. A Stochastic Petri Net (SPN) is a pair $S = (N, \mu)$ where N is a PN and for all $t \in T, \mu(t) \in \mathbb{R}^+$ is the rate of firing of transition t .

In a given marking, a delay is computed for every enabled transition t with an exponential probability distribution function of parameter $\mu(t)$. A SPN system has a time semantic [18, 15] that is defined according to: (a) a *single server policy*: each transition can only be fired once by a given marking; (b) a *race policy*: the transition whose firing delay elapses first is assumed to be the one that will fire next; (c) a *resampling memory policy*: at the entrance in a marking, the remaining delays associated with all transitions are forgotten.

Similarly to a PN, a *Stochastic Net System* (SNS) is a pair $\langle S, M_0 \rangle$ where S is an SPN and M_0 is an initial marking. Definitions for NS are transposed to SNS. Given a sequence $\sigma \in T^*$, we write $C(\sigma)$ for the set of infinite sequences prefixed by σ , i.e. $C(\sigma) = \{\sigma' \in T^\omega \mid \exists \sigma'' \in T^\omega : \sigma' = \sigma\sigma''\}$. The set of infinite sequences is the support of a probability measure defined by Caratheodory's extension theorem from the probabilities of the cylinders: the probability of the cylinder starting by the empty sequence ε is equal to 1 and, for σt a sequence, the probability of $C(\sigma t)$ in M_0 , written $\mathbb{P}(\sigma t)$, satisfies

$$\mathbb{P}(\sigma t) = \mathbb{P}(\sigma) \times \frac{\mu(t)}{\sum_{t' \in T, M_0[\sigma t']} \mu(t')}$$

In the following, we want to use the previous definitions to deal with the problem of fault diagnosis where the goal is to detect the occurrence of a fault under partial observation. To this aim, we associate a well precise physical meaning to implicit, explicit, and relevant transitions. In more detail:

- Implicit transitions correspond to transitions that cannot be observed. They are called *silent* or *unobservable* and could either model a regular (nominal) behaviour or a faulty behaviour of the system.
- Conversely, explicit transitions model transitions that can be observed. Those *observable* transitions are assumed to be a regular behaviour of the system
- The set of faulty transitions is chosen as the set of relevant transitions.

We denote the above three sets as T_u , T_o , and T_f , respectively and choose $T_e = T_o$ and $T_i = T_u$.

In simple words, we may assume that observable transitions model events whose occurrence is detected by the presence of a sensor. On the contrary, unobservable transitions correspond to events to whom no sensor is associated. Note that, in the general case, the same output signal may correspond to different events (different transition firings). This can be easily modelled using the notion of *labelling function*. $\mathcal{L} : T \rightarrow L \cup \{\varepsilon\}$ that assigns to each transition $t \in T$ either a symbol from a given alphabet of events L (if $T \in T_o$) or the empty string ε (if $T \in T_u$). We extend naturally \mathcal{L} to sequences of transitions with $\mathcal{L}(\sigma t) = \mathcal{L}(\sigma)\mathcal{L}(t)$. The observed word w of events associated with the sequence σ is $w = \mathcal{L}(\sigma)$. Note that the length of a sequence σ is always greater than or equal to the length of the corresponding word w (denoted $|w|$). In fact, if σ contains k' transitions in T_u then $|\sigma| = k' + |w|$. Given a word $w \in L^*$, we write $\mathbb{P}(w) = \sum_{\sigma \in P_e^{-1}(w)} \mathbb{P}(\sigma)$. Assuming (A2), this sum is finite.

Example 5. Consider again the NS in Figure 1, where the labelling function \mathcal{L} is such that $\mathcal{L}(t_1) = b$, $\mathcal{L}(t_2) = a$, $\mathcal{L}(t_3) = \mathcal{L}(t_4) = \varepsilon$ and $\mathcal{L}(t_5) = \mathcal{L}(t_6) = c$. Thus, t_3 and t_4 are unobservable.

Transition t_5 being observable, the T_u -induced subnet is acyclic.

The goal of diagnosis is to detect whether a faulty event occurred in the system. We denote by $T_f \subseteq T_u$ the set of faulty transitions. A sequence σ is faulty

if there exists $t \in T_f$ such that $t \in \sigma$, otherwise it is correct. An observed word w is surely faulty (resp. correct) iff every sequence σ with $\mathcal{L}(\sigma) = w$ is faulty (resp. correct) sequences, otherwise it is ambiguous. An NS system is diagnosable iff all faults can be detected after a finite delay.

Definition 10. *An NS $\langle N, M_0 \rangle$ is diagnosable if for every faulty sequence σ enabled by M_0 , there exists $n \in \mathbb{N}$ such that for all sequences $\sigma' \in T^n$ with $\sigma\sigma'$ enabled by M_0 , $\mathcal{L}(\sigma\sigma')$ is surely faulty.*

A similar notion of diagnosability (called FF-diagnosability in [2], A-diagnosability in [26]) can be defined for SNS. In simple words, faults need not to be detected for sure, but need to be detected almost surely.

Definition 11. *An NS $\langle N, M_0 \rangle$ is FF-diagnosable if for every faulty sequence σ enabled by M_0 , we have*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\{\sigma' \in T^n \mid \mathcal{L}(\sigma\sigma') \text{ is not surely faulty}\}) = 0.$$

Example 6. Consider the NS in Figure 1, with $\mathcal{L}(t_5) = \mathcal{L}(t_6) = a$, $T_u = \{t_3, t_4\}$ and $T_f = \{t_3\}$. The sequence $\sigma_f = t_1 t_2 t_3 (t_1)^\omega$ is faulty but its observed word is ambiguous, thus this NS is not diagnosable. However any sequence containing more than two a is surely faulty and, adding a rate $\mu(t) = 1$ to every transition t , with probability 1 a faulty sequence will trigger t_5 infinitely often. Therefore the associated SNS is FF-diagnosable.

4 Diagnosability Analysis of Stochastic Bounded Net Systems

Diagnosability analysis is known to be EXPSPACE-complete for bounded NS. Using the basis reachability graph, the authors of [8] gave an algorithm which, although still EXPSPACE, is far more efficient than the previous ones. Similarly, since FF-diagnosability is PSPACE-complete for Markov chains and one could transform a bounded stochastic Petri net into a Markov chain exponential in the size of the net (in the number of places, transitions and on the maximum number of tokens in the net), FF-diagnosability is in EXPSPACE. Moreover, the proof of EXPSPACE-hardness of diagnosability from [1] can be directly used for FF-diagnosability as the Petri net they build is diagnosable iff it is FF-diagnosable. Thus we can state the following result.

Theorem 1. *The FF-diagnosability analysis is EXPSPACE-hard.*

As for diagnosability analysis, the BCG can be used to reduce the computation cost. The system being bounded here, the BCG reduces in fact to the basis reachability graph. The rest of this section will be devoted to explaining how to use the BCG to analyse the FF-diagnosability of an NS.

Our first step is to define the *belief automaton* [14] associated with a BCG. The state of the belief automaton, called belief, reached after an observation

w contains the set of basis markings reachable with a sequence of observations w . Moreover, those markings are paired with a tag expressing the following properties:

- F tags the basis markings which were reached using a faulty transition,
- C marks the others.

The belief automaton is deterministic and exponential in the size of the BCG. It is similar to a form of determinisation of the nondeterministic automaton obtained from the reachability graph labeling the arcs with transition labels (as opposed to labeling the arcs with transitions) which, in the context of Discrete Event Systems, is called "observer" [10].

Definition 12. Let $\langle S, M_0 \rangle$ be an SNS with set of unobservable transitions T_u and let $BCG_{\langle N, M_0 \rangle}^{T_f} = (\mathcal{M}, M_0, \Delta)$ be its BCG with relevant set of transitions T_f . The belief automaton $B_G = \{Q_B, \{(M_0, C)\}, \Delta_B\}$ associated with $BCG_{\langle N, M_0 \rangle}^{T_f}$ is defined by:

- $Q_B = 2^{\mathcal{M} \times \{F, C\}}$;
- $(B, a, B') \in \Delta_B$ where:
 - $(M', F) \in B'$ iff there exists $(M, C) \in B$, a transition $(p, t) \in \Delta$ with $\mathcal{L}(t) = a$ and p is the firing vector of a faulty sequence, or there exists $(M, F) \in B$, a transition $(p, t) \in \Delta$ with $\mathcal{L}(t) = a$;
 - $(M', C) \in B'$ iff there exists $(M, C) \in B$, a transition $(p, t) \in \Delta$ with $\mathcal{L}(t) = a$ and p is the firing vector of a correct sequence.

A maximally strongly connected component (SCC) \mathcal{C} of the belief automaton associated with $BCG_{\langle N, M_0 \rangle}^{T_f}$ B_G is called *terminal ambiguous SCC* if there exists a belief B of \mathcal{C} , two markings $(M_B, F), (M'_B, C) \in B$ and a marking $M \in R_i(N, M_B)$ such that for every sequence σ enabled by M , for B' the belief of B_G such that $B[\mathcal{L}(\sigma)]B'$, $B' \in \mathcal{C}$. In other words, \mathcal{C} is "ambiguous" (i.e. contains markings tagged by F and markings tagged by C) and there is a marking associated with B which, once reached, implies that the rest of the run will have its belief remain in \mathcal{C} . Such a pair (M, B) is called a witness of the terminality of \mathcal{C} .

We can now characterise FF-diagnosability based on the belief automaton.

Lemma 1. Let $\langle S, M_0 \rangle$ be an SNS with set of unobservable transitions T_u verifying Assumptions (A1) and (A2), $\langle S, M_0 \rangle$ is FF-diagnosable iff the belief automaton associated with $BCG_{\langle N, M_0 \rangle}^{T_f}$ does not contain any terminal ambiguous SCC.

Proof. Suppose there exists a terminal ambiguous SCC \mathcal{C} with witness (M, B) . Let $\sigma = \sigma_1 t_1 \dots \sigma_n t_n$ be a faulty sequence in the BCG such that $M_0[\sigma]M$ and $\mathcal{L}(\sigma)$ leads to B in B_G . Let σ' be a sequence such that $\sigma\sigma'$ is enabled by M_0 . Let B' be the state of the belief automaton reached from B by observing $\mathcal{L}(\sigma')$, then B' contains an element of the form (M', C) as B is reachable from

B' due to the terminality of the SCC. Therefore, by definition of the belief automaton, there exists a correct sequence $\tilde{\sigma}$ such that $\mathcal{L}(\tilde{\sigma}) = \mathcal{L}(\sigma\sigma')$. Thus $\mathcal{L}(\sigma\sigma')$ is not surely faulty. As this is true for every sequence extending σ , we have $\lim_{n \rightarrow \infty} \mathbb{P}(\{\sigma' \in T^n \mid \mathcal{L}(\sigma\sigma') \text{ is not surely faulty}\}) \geq \mathbb{P}(\sigma) > 0$ which implies that the SNS is not FF-diagnosable.

Conversely, suppose that there is no terminal ambiguous SCC. Let $\sigma = \sigma_1 t_1 \dots \sigma_n t_n$ be a faulty sequence of the SNS with σ_i sequences of unobservable transitions and t_i observable. Let M be the marking such that $M_0[\sigma]M$, B the belief reached in B_G by observing $\mathcal{L}(\sigma)$ and \mathcal{C} the maximal SCC B belongs to. B_G contains a marking (M_B, F) as M is reached by a faulty sequence which fault is prior to the last observation. Indeed, supposing for simplicity (possible because the NS verifies (A1)) that the σ_i are minimal explanation, then as σ is faulty, one of them belongs to $\Sigma_{min}^{T_f}$. If B does not contain a marking (M'_B, C) then $\mathcal{L}(\sigma)$ is surely faulty. Else, as \mathcal{C} is not terminal ambiguous, for any marking \hat{M} reached by a sequence which observation ends in \mathcal{C} , there is a sequence $\sigma_{\hat{M}}$ such that observing $\mathcal{L}(\sigma_{\hat{M}})$ exits \mathcal{C} . As there is a finite number of pairs of marking and belief due to the NS being bounded, the minimum probability of such a sequence has a non null lower bound. Therefore, the set of sequences extending σ and which observation stays in \mathcal{C} has probability 0. In other words, any maximal SCC in which a positive measure of sequences extending σ stays infinitely only contains belief with no marking tagged by C . Thus those sequences are surely faulty. Hence $\lim_{n \rightarrow \infty} \mathbb{P}(\{\sigma' \in T^n \mid \mathcal{L}(\sigma\sigma') \text{ is not surely faulty}\}) = 0$, the PN is FF-diagnosable. \square

Theorem 2. *Given an SNS $\langle S, M_0 \rangle$ with set of unobservable transitions T_u bounded by a value $k \in \mathbb{N}$ and satisfying Assumptions (A1) and (A2), FF-diagnosability analysis is in EXPSPACE (in the size of the PN and in k) using the BCG with relevant set of transitions T_f .*

Proof. We will explain now how to check the characterisation given by Lemma 1 is in EXPSPACE in the following way:

- We first guess a belief B containing at least one marking tagged by C and one marking tagged by F and a marking M . We will now check if this pair (M, B) is a witness of the terminality of an SCC.
- We verify that M is reachable by unobservable transitions from a basis marking of B tagged by F by guessing the unobservable path.
- We will now verify that the SCC B belongs to is terminal once M is reached. This is done by guessing a belief B' that would be outside of this SCC and a (at most doubly exponential) sequence σ enabled by M such that $\mathcal{L}(\sigma)$ leads from B to B' . We then verify that this belief B' is indeed outside the SCC by guessing a sequence σ' such that $\sigma\sigma'$ is enabled by M and $\mathcal{L}(\sigma')$ leads from B' to B . If σ' can be found, B' was a wrong guess and therefore B belongs indeed to a terminal SCC.

Every guess is removed using the theorem of Savitch [25] at every step. The algorithm is EXPSPACE as it only needs to keep one belief in memory while

visiting the graph with a sequence which is at most doubly exponential. In practice, one could also build the full graph which would allow to remove the guesses, but doing so would technically raise the complexity to 2-EXPTIME.

As B can be checked to be part of a terminal ambiguous SCC in EXPSPACE and its reachability in B_G can also be verified in EXPSPACE, the given algorithm is EXPSPACE. Therefore according to the characterisation of Lemma 1, the FF-diagnosability analysis can be done in EXPSPACE. \square

The notion of terminal ambiguity could be removed by making the product of B_G with the reachability graph and checking the ambiguity of the faulty ambiguous bottom SCC (SCC from which no other SCC is reachable) of the product. This would be closer to the techniques used for Markov chains for example in [2] but would require the construction of the entire reachability graph (although we would still retain most of the efficiency as only the BCG is used for the exponential construction of B_G).

5 Diagnosability Analysis of Unbounded Net Systems

We now focus on unbounded systems for which the BCG was developed. Unfortunately, the basis coverability graph can not be used to decide FF-diagnosability. In fact, we will show here that for unbounded stochastic Petri nets, FF-diagnosability is undecidable. To do so we will reduce the problem of the language inclusion for Petri nets, namely: given two NS $\langle N^1, M_0^1 \rangle$ and $\langle N^2, M_0^2 \rangle$ does $\mathcal{L}(L(N^1, M_0^1)) \subseteq \mathcal{L}(L(N^2, M_0^2))$ hold? This problem is known to be undecidable [11].

Theorem 3. *The FF-diagnosability analysis of unbounded SNS is undecidable.*

Proof. See Appendix.

FF-diagnosability was also shown undecidable for probabilistic pushdown automata [3] which are another probabilistic model representing infinite state systems. However, there is a known restriction for which the problem becomes decidable [3]: probabilistic visibly pushdown automata for which the set of observations is divided into those corresponding to an action adding an element to the stack, those removing one and those that do not modify the size of the stack, unobservable actions can not modify the size of the stack. This way, an observer knows at all time the size (but not the content) of the stack. Mimicking this restriction for unbounded NS would require that the labelling function allows to know at all time how many tokens are in the system. This is the case in the reduction realised in the proof of the previous result however, thus such a restriction for unbounded NS would remain undecidable.

Diagnosability on the contrary was proven decidable [5, 1]. To do so, the authors of [5] gave a characterisation of diagnosability using a tool called Verifier Net. The verifier net is obtained by a composition (related to a parallel composition of the studied NS and its $T \setminus T_f$ -induced subnet with synchronisation on the observable transitions.

Definition 13. Given an NS $\langle N, M_0 \rangle$, let $\langle N', M'_0 \rangle$ be the $T \setminus T_f$ -induced subnet of $\langle N, M_0 \rangle$ (prime are used to differentiate states and transitions of N' from those of N). We build the verifier net (VN) $\langle \tilde{N}, \tilde{M}_0 \rangle$ of $\langle N, M_0 \rangle$ with $\tilde{N} = (\tilde{P}, \tilde{T}, \tilde{Pre}, \tilde{Post})$ where:

- $\tilde{P} = P \cup P'$,
 - $\tilde{T} = (T'_o \times T_o) \cup (T \setminus T_f \times \{\lambda\}) \cup (\{\lambda\} \times T)$,
 - for $t \in T, t' \in T' \setminus T_f, p \in P$, and $p' \in P'$, we have
 - $\tilde{Pre}(p, (\lambda, t)) = Pre(p, t)$ and $\tilde{Post}(p, (\lambda, t)) = Post(p, t)$,
 - $\tilde{Pre}(p', (t', \lambda)) = Pre(p', t')$ and $\tilde{Post}(p', (t', \lambda)) = Post(p', t')$,
 - if $\mathcal{L}(t) = \mathcal{L}(t') \neq \varepsilon$, $\tilde{Pre}(p', (t', t)) = Pre(p', t')$ and $\tilde{Post}(p', (t', t)) = Post(p', t')$, $\tilde{Pre}(p, (t', t)) = Pre(p, t)$ and $\tilde{Post}(p, (t', t)) = Post(p, t)$.
- All unspecified values are equal to 0.

Theorem 4 ([5]). An NS $\langle N, M_0 \rangle$ verifying Assumption (A1) is diagnosable iff there does not exist any cycle in the coverability graph of the VN which (1) starts from an ω -marking reachable by a faulty sequence and (2) is associated with a repetitive sequence in the associated VN.

We will now use this characterisation to formulate a similar one using the BCG instead of the coverability graph. A sequence of the BCG is called faulty if one of the minimal explanations used belong to $\Sigma_{min}^{T_f}$.

Theorem 5. An NS $\langle P, M_0 \rangle$ verifying Assumptions (A1) and (A2) is diagnosable iff there does not exist any cycle in the BCG with relevant set of transitions T_f of the VN which (1) starts from a basis marking reachable by a faulty sequence and (2) is associated with a repetitive sequence in the associated VN.

Proof. We will show that the existence of such a cycle in the BCG is equivalent to the existence of this cycle in the coverability graph.

Supposing there exists a cycle associated with a firable repetitive sequence $\sigma \in T^*$ in the associated VN that starts from a basis marking M_ω reached by a faulty sequence in the BCG with relevant set of transition T_f of the VN, then by Proposition 4, M_ω is an ω -marking of the coverability graph and by construction of the BCG, there exists a directed cycle starting in M_ω in the coverability graph whose arcs form σ .

Now suppose that there is a firable repetitive sequence $\sigma = \sigma_1 t_1 \dots \sigma_n t_n$ in the VN that is associated to a cycle starting from an ω -marking reached by a faulty sequence in the coverability graph of the VN. There thus exists a marking M of the VN such that σ is repetitive starting in M . Because of the assumption (A2), σ contains at least one observable transition. According to Proposition 9, there thus exists a basis marking M_ω and an ω -marking M_u such that $M_u \in R_i(N, M_\omega)$, $M_u \geq_\omega M$ and there is a $k \in \mathbb{N}$ and a directed cycle starting in M_ω whose arcs, projected on the second component, form $P_o(\sigma)^k$. Moreover, as M is reached by a faulty sequence $\sigma' = \sigma'_1 t'_1 \dots \sigma'_n t'_n \sigma'_{n+1}$, one can choose M_ω to be reached by a sequence that used a minimal explanation from $\Sigma_{min}^{T_f}$: if σ'_i is faulty, one can choose the minimal explanation of t_i to belong in $\Sigma_{min}^{T_f}$.

Consequently the characterisation of Theorem 4 and Theorem 5 are equivalent and can both be used to solve diagnosability. \square

Conclusion

In this paper, we introduced the notion of basis coverability graph which provides an abstracted representation of the coverability graph. We established multiple properties of the basis coverability graph, especially how it can be used to approximate the reachability set efficiently. We then focused on diagnosability and stochastic diagnosability, showed how the basis reachability graph can be employed to solve some of those problems and showed undecidability when it can not be used. The logical next step would be to implement the algorithms obtained and compare their efficiency with other algorithms ([1] for example) on case studies.

References

1. B. Bérard, S. Haar, S. Schmitz, and S. Schwoon. The Complexity of Diagnosability and Opacity Verification for Petri Nets. In *Petri nets 2017*, Lecture Notes in Computer Science. Springer, 2017.
2. N. Bertrand, S. Haddad, and E. Lefaucheux. Foundation of diagnosis and predictability in probabilistic systems. In *Proceedings of FSTTCS'14*, volume 29 of *LIPICs*, pages 417–429. Leibniz-Zentrum für Informatik, 2014.
3. N. Bertrand, S. Haddad, and E. Lefaucheux. Diagnosis in infinite-state probabilistic systems. In *Proceedings of CONCUR'16*, volume 59 of *LIPICs*, pages 37:1–37:14. Leibniz-Zentrum für Informatik, 2016.
4. H. Boucheneb and K. Barkaoui. Reducing interleaving semantics redundancy in reachability analysis of time Petri nets. *ACM Trans. Embed. Comput. Syst.*, 12(1):7:1–7:24, January 2013.
5. M. P. Cabasino, A. Giua, S. Lafortune, and C. Seatzu. A new approach for diagnosability analysis of petri nets using verifier nets. *IEEE Transactions on Automatic Control*, 57(12):3104–3117, Dec 2012.
6. M. P. Cabasino, A. Giua, and C. Seatzu. Fault detection for discrete event systems using Petri nets with unobservable transitions. *Automatica*, 46(9):1531–1539, 2010.
7. M. P. Cabasino, A. Giua, and C. Seatzu. *Introduction to Petri Nets*, pages 191–211. Springer London, 2013.
8. M. P. Cabasino, A. Giua, and C. Seatzu. Diagnosability of discrete-event systems using labeled Petri nets. *IEEE Trans. Automation Science and Engineering*, 11(1):144–153, 2014.
9. M.P. Cabasino, A. Giua, M. Poggi, and C. Seatzu. Discrete event diagnosis using labeled Petri nets. an application to manufacturing systems. *Control Engineering Practice*, 19(9):989 – 1001, 2011.
10. C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems - Second Edition*. Springer, 2008.
11. J. Esparza and M. Nielsen. Decidability issues for Petri nets - a survey, 1994.
12. A. Giua, C. Seatzu, and D. Corona. Marking estimation of Petri nets with silent transitions. *IEEE Transactions on Automatic Control*, 52(9):1695–1699, Sept 2007.
13. P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem*, volume 1032. Springer, 1996.
14. S. Haar, S. Haddad, T. Melliti, and S. Schwoon. Optimal constructions for active diagnosis. In *Proceedings of FSTTCS'13*, volume 24 of *LIPICs*, pages 527–539. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.

15. S. Haddad and P. Moreaux. *Stochastic Petri Nets*. Wiley-ISTE, 2009.
16. R. M. Karp and R. E. Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.
17. Z.Y. Ma, Y. Tong, Z.W. Li, and A. Giua. Basis marking representation of Petri net reachability spaces and its application to the reachability problem. *IEEE Transactions on Automatic Control*, 62(3):1078–1093, 2017.
18. M. A. Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets*. John Wiley & Sons, Inc., 1994.
19. E. W. Mayr. An algorithm for the general Petri net reachability problem. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*, STOC '81, pages 238–246. ACM, 1981.
20. T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.
21. M. Nielsen, G. Plotkin, and G. Winskel. Petri nets, event structures and domains, part I. *Theoretical Computer Science*, 13(1):85 – 108, 1981. Special Issue Semantics of Concurrent Computation.
22. P.-A. Reynier and F. Servais. Minimal coverability set for Petri nets: Karp and miller algorithm with pruning. *Fundamenta Informaticae*, 122(1-2):1–30, January 2013.
23. R.Lipton. The Reachability Problem Requires Exponential Space. Technical report, Yale University, 1976.
24. M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE Trans. Aut. Cont.*, 40(9):1555–1575, 1995.
25. W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177 – 192, 1970.
26. D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Transactions on Automatic Control*, 50(4):476–492, 2005.
27. Y. Tong, Z. Li, C. Seatzu, and A. Giua. Verification of state-based opacity using Petri nets. *IEEE Transactions on Automatic Control*, 62(6):2823–2837, June 2017.
28. A. Valmari. *The state explosion problem*, pages 429–528. Springer Berlin Heidelberg, 1998.

A Appendix

This appendix contains the proof of undecidability of the FF-diagnosability analysis for unbounded Petri nets in order to help the reviewing process. it will be omitted in the final version, however a HAL link to this proof will be given.

Proof. Let $\langle N^1, M_0^1 \rangle$ and $\langle N^2, M_0^2 \rangle$ be two PN over alphabet Σ . For simplicity we will suppose the initial marking M_0^i to have a single token on a place p_0^i for $i = 1, 2$, that every transition is observable and that the number of tokens in the system to be equal to the length of the sequence plus 1 which can be done without loss of generality.

We build the SN $(\langle N, M_0 \rangle, \mu)$ (represented in Figure 4) where:

- $P = P^1 \cup P^2 \cup \{p_0\} \cup \{p_{emp}^i, p_{run}^i, p_{err}^i \mid i = 1, 2\}$;
- $T = T^1 \cup T^2 \cup \{t_{in}^i, t_{\#}^i, t_{rese}^i, t_{resn}^i \mid i = 1, 2\} \cup \{t_a^i \mid a \in \Sigma, i = 1, 2\} \cup \{t_{emp}^p, t_{err}^p \mid p \in P^1 \cup P^2\}$
- for $i \in \{1, 2\}, p \in P^i, t \in T^i, Pre(p, t) = Pre^i(p, t)$ and $Pre(p_{run}^i, t) = 1, Pre(p_0, t_{in}^i) = 1, Pre(p_{emp}^i, t_{\#}^i) = 1, Pre(p_{run}^i, t_{resn}^i) = 1, Pre(p_{err}^i, t_{rese}^i) = 1, \text{ for } a \in \Sigma, Pre(p_{err}^i, t_a^i) = 1, Pre(p, t_{emp}^p) = Pre(p_{emp}^i, t_{emp}^p) = 1, Pre(p, t_{err}^p) = Pre(p_{emp}^i, t_{err}^p) = 1$. When undefined, $Pre(p, t) = 0$.
- for $i \in \{1, 2\}, p \in P^i, t \in T^i, Post(p, t) = Post^i(p, t)$ and $Post(p_{run}^i, t) = 1, Post(p_0, t_{in}^i) = Post(p_{run}^i, t_{in}^i) = 1, Post(p_{run}^i, t_{\#}^i) = Post(p_0, t_{\#}^i) = 1, Post(p_{emp}^i, t_{resn}^i), Post(p_{emp}^i, t_{rese}^i), \text{ for } a \in \Sigma, Post(p_{err}^i, t_a^i) = Post(p_0, t_a^i) = 1, Post(p_{emp}^p, t_{emp}^p) = 1, Post(p_0, t_{err}^p) = 2, Post(p_{err}^p, t_{err}^p) = 1$. When undefined, $Post(p, t) = 0$.
- We suppose the observation function, \mathcal{L} , defined on N^1 and N^2 we extend it on N by, for $p \in P^1 \cup P^2, i \in \{1, 2\}, a \in \Sigma, \mathcal{L}(t_{in}^i) = \varepsilon, \mathcal{L}(t_{\#}^i) = \mathcal{L}(t_{err}^p) = \mathcal{L}(t_{rese}^i) = \mathcal{L}(t_{resn}^i) = \#, \mathcal{L}(t_a^i) = a, \mathcal{L}(t_{emp}^p) = \flat$.
- for $i \in \{1, 2\}, p \in P_i, \mu(t_{rese}^1) = \mu(t_{resn}^1) = \mu(t_{emp}^p) = 2|T_1|$ (assuming $|T_1| \geq 1$) and for every other transition $t \mu(t) = 1$.

Moreover, t_{in}^1 is a faulty transition.




Fig. 4. Reduction from language inclusion. The Figure 5 represents the content of the box Box^1 , it is similar for Box^2 .




Fig. 5. Content of the box Box^1 .

Informally, on the first transition the system randomly activates one of the two box represented in the Figure 4, the first one being reached by a faulty transition. If we reach the box i , a word $w \in \mathcal{L}(L(N^i, M_0^i))$ is observed followed by a \sharp then we observe a certain number of \flat followed by another \sharp . If the number of \flat is equal to the length of w , then we repeat the operation as a new word of $\mathcal{L}(L(N^i, M_0^i))$ is observed. If there is less \flat (there can not be more), then there are leftover tokens in the net when the second \sharp is fired. This allows the net to read any word $w \in \Sigma^*$ before starting to empty the net again. In other words, we learn informations on the system iff it emptied itself correctly before a \sharp is read. We will show here that the system is FF-diagnosable iff $\mathcal{L}(L(N^1, M_0^1)) \not\subseteq \mathcal{L}(L(N^2, M_0^2))$.

First remark that the set of observed words of the infinite sequences of the SNS starting by the transition t_{in}^i , denoted L^i , contains exactly the words of the form $w_1 \sharp \flat^{n_1} \sharp \dots w_k \sharp \flat^{n_k} \sharp \dots$ where for all $1 \leq j \leq k$, (1) $w_j \in \Sigma^*$, (2) $\sum_{m=1}^j |w_m| + 1 \geq \sum_{m=1}^j n_m$ and (3) $w_j \in \mathcal{L}(L(N^i, M_0^i))$ if $\sum_{m=1}^{j-1} |w_m| + 1 = \sum_{m=1}^{j-1} n_m$.

Suppose that $\mathcal{L}(L(N^1, M_0^1)) \subseteq \mathcal{L}(L(N^2, M_0^2))$. Let σ be a finite faulty sequence. As σ is faulty, it initially fired t_{in}^1 , thus $\mathcal{L}(\sigma) \in L^1$. Thanks to the above remark on the languages L_i , and as $\mathcal{L}(L(N^1, M_0^1)) \subseteq \mathcal{L}(L(N^2, M_0^2))$, $\mathcal{L}(\sigma) \in L^2$, therefore there exists a sequence σ' starting by the transition t_{in}^2 with same observation as σ . Moreover this transition is not faulty as it did not fire t_{in}^1 initially and can not fire it after the first transition. Therefore $\mathcal{L}(\sigma)$ is not surely faulty. As this is true for every faulty sequence, the system is not FF-diagnosable.

Suppose now that $\mathcal{L}(L(N^1, M_0^1)) \not\subseteq \mathcal{L}(L(N^2, M_0^2))$. There thus exists a word w such that $w \in \mathcal{L}(L(N^1, M_0^1)) \setminus \mathcal{L}(L(N^2, M_0^2))$. The observed words of L_1 of

the form $w_1 \#^{n_1} \# \dots w_k \#$ where $\sum_{m=1}^{k-1} |w_m| = \sum_{m=1}^{k-1} n_m$ and $w_k = w$ are surely faulty as they do not belong into L^2 . We denote SL_1 the set of those observed words. Let us show that with probability 1 an infinite faulty sequence is prefixed by a sequence whose observation belongs to SL_1 .

Due to the choice of the rates μ , the system is more likely to remove a token than to add one in the PN N_1 . Therefore, with probability 1, a faulty sequence will infinitely often trigger $t_{\#}^1$ while there is no token in P^1 . Therefore with probability 1, the observation of a faulty sequence will be of the form $w_1 \#^{n_1} \# \dots w_k \#^{n_k} \# \dots \in L^1$ with infinitely many $j \in \mathbb{N}$ such that $\sum_{m=1}^{j-1} |w_m| + 1 = \sum_{m=1}^{j-1} n_m$. There is a probability $p > 0$ that for any such j , $w_j = w$ as $w \in \mathcal{L}(L(N^1, M_0^1))$. Therefore with probability 1, there exists $j \in \mathbb{N}$ such that $w_j = w$. Hence with probability 1 an infinite faulty sequence will have a prefix whose observation belongs to SL_1 . This implies that the SNS is FF-diagnosable. \square