Structural characterization and qualitative properties of Product Form Stochastic Petri Nets*

S. Haddad[1], P. Moreaux[1,2], M. Sereno[3], and M. Silva[4]

[1] LAMSADE, Université Paris Dauphine,
{haddad,moreaux}@lamsade.dauphine.fr,
[2] LERI-RESYCOM, Université de Reims Champagne-Ardenne,
[3] Dipartimento di Informatica, Università di Torino,
matteo@di.unito.it,
[4] Dep. de Ingenieria Elec.a e Informatica, Universidad de Zaragoza,
silva@posta.unizar.es

Abstract. The model of Stochastic Petri nets (SPN) with a product form solution (Π-net) is a class of nets for which there is an analytic expression of the steady state probabilities w.r.t. markings, as for product form queueing networks w.r.t. queue lengths. In this paper, we prove new important properties of this kind of nets. First we provide a polynomial time (w.r.t. the size of the net structure) algorithm to check whether a SPN is a Π-net. Then, we give a purely structural characterization of SPN for which a product form solution exists regardless the particular values of probabilistic parameters of the SPN. We call such nets $\overline{\Pi}$-nets. We also present untimed properties of Π-nets and $\overline{\Pi}$-nets such like liveness, reachability, deadlock freeness and characterization of reachable markings. The complexity of the reachability and the liveness problems is also addressed for Π-nets and $\overline{\Pi}$-nets. These results complement previous studies on these classes of nets and improve the applicability of Product Form solutions.

1 Introduction

Stochastic Petri nets (SPNs) are a powerful tool for modelling and evaluating the performance systems involving concurrency, non determinism, and synchronization, such as parallel and distributed systems, communication networks, etc. The stochastic semantics of SPN have been proven to be a Continuous Time Markov Chain and steady state analysis can thus be expressed as the solution of a system of equilibrium equations, one for each possible marking of their state space. The major problem in the computation of performance measures using SPNs is thus the size of the reachability set of these models that increases exponentially both with the number of tokens in the initial marking and with the number of places in the net. As a consequence, the dimension of this reachability set and the time complexity of the solution procedure preclude, in the general

* At time of writing, P. Moreaux was visiting professor at the Dipartimento di Informatica, Università di Torino

case, the direct exact numerical evaluation of many interesting models. To cope with this problem, we can first accept non exact performance measures. The two main approaches developed in this area are discrete-event simulation and approximate methods. Bounds computation methods provide more reliable information about the the performance indices. However, if we wish to obtain the exact values of performance measures, then we may improve numerical methods solving the underlying mathematical problem (linear or differential systems of equations) and/or we may relate the structure of the model to the properties of these underlying mathematical objects.

One successful approach in this last direction is the product form analysis (PFA) for Queueing Networks (QN), that is the expression of basic performance indices of QN, such as steady state probabilities, mean throughputs, utilization, etc., as functions of the model parameters (service rates, routing probabilities, properties of the service stations, etc.). The first structural property involved in PFA is obviously the setting up of the model as a collection of service stations bounded with paths taken by "clients". From this structure, PF solutions may be proven for several classes of QN by examination of sets of some kind of "local balance equations", for instance equations established for each station. Second, specific descriptions of the state space of PF-QN lead to important relations. For instance, the convolution algorithms [21] and the Mean Value Analysis (MVA) method [4] are based upon recursive relations between models with state spaces with different number of clients. Unfortunately, (the standard version of) PF-QN offer limited possibilities for what concerns synchronization between clients activities. This situation was one of the main motivations in the study of Stochastic Petri Nets (SPN) with a Product form solution (PF-SPN). First results about PF-SPN were established in [15] based on the structure of the reachability graph of the net. Recently, several authors proposed structural sufficient conditions for a Petri net to be a PF-SPN. These results are summarized in Section 2. The present paper supplements previous results for PF-SPN regarding four important issues.

Membership Problem for SPN with PF solution As we will see in Section 2, a straightforward verification procedure for deciding whether a given SPN has a PF solution requires the computation of all minimal T-semiflows of the marked net (T-semiflows are structural invariants of Petri nets (PN), see Section 2). It is however known that the number of minimal T-semiflows can be exponential in the number of transitions (e.g., [17]). In fact, we establish a polynomial time algorithm to decide whether a SPN has a PF solution.

Rate independent structural characterization of PF-SPN Previous criteria for PF-SPN have two drawbacks: they are only sufficient conditions, and they involve properties of the rates of the transitions of the net. We present a necessary and sufficient structural condition on nets to admit a PF solution whatever the rates of its transitions. Hence we prove a rate-independent structural characterization of PF-SPN. Moreover, this criterion can be checked in polynomial time.

Untimed properties of PF-SPN We investigate untimed properties for the class of PF-SPN. Since many results (deadlock-freeness, liveness, etc.) have been established for several known classes of PN, it can be valuable to point out the relation between PF-SPN and these classes.

Reachability Set properties Efficient numerical solutions for PF-SPN} require to characterize subsets of reachable markings. It is hence important to have a structural criterion for reachable markings (e.g., a method based on the minimal P-semiflows, a method based on the net state equation, etc.). We present new results about these possible criteria.

The organization of the paper is the following: in Section 2, we review SPN and previous results about Π-nets. Section 3 presents the verification procedure for PF-SPN and a series of results about the class of PF-SPN in relation to other classes of Petri nets. In Section 3.3 we define the new class, $\overline{\Pi}$-nets, of PF-SPN corresponding to rate independent criteria for a PF solution together with globally dependent rates. Untimed properties of Π-nets are studied in Section 4. The conclusion summarizes results presented in the paper.

2 Background and notations

2.1 Stochastic Petri nets

One may find introductory presentations of Petri net concepts for instance in [19, 20, 26]. We remind the reader only with definitions necessary to understand product form results for stochastic Petri nets.

A *marked stochastic Petri net* is a 5-tuple $SPN = (\mathcal{P}, \mathcal{T}, W, \mathcal{Q}, \mathbf{m_0})$, where $\mathcal{P}$ and $\mathcal{T}$ are disjoint sets of *places* and *transitions* (with $|\mathcal{P}| = np$ and $|\mathcal{T}| = nt$), $W := (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P}) \to \mathbb{N}$ defines the *weighted flow relation*: if $W(j,i) > 0$ (resp. $W(i,j) > 0$) then we say that there is an *arc* from t_j to p_i, with *weight* or *multiplicity* $W(j,i)$ (resp. there is an arc from p_i to t_j with weight $W(i,j)$), $\mathcal{Q}$ is the set of transition firing rates drawn from exponential distributions, and $\mathbf{m_0}$ is the initial marking.

For a given transition $t_j \in \mathcal{T}$, its *preset* and *postset* are given by ${}^\bullet t_j = \{p_i \mid W(i,j) > 0\}$ and $t_j{}^\bullet = \{p_i \mid W(j,i) > 0\}$, respectively. In the same manner we can define the *preset* and *postset* of a given place.

For any transition t_j, from the weighted flow relation we can the define the *input vector* $\mathbf{i}(t_j) = [W(1,j), W(2,j), \dots, W(|\mathcal{P}|, j)]$ and the *output vector* $\mathbf{o}(t_j) = [W(j,1), W(j,2), \dots, W(j, |\mathcal{P}|)]$. From the weighted flow relation we can also define the *incidence matrix* $\mathbf{C}$ with entries $\mathbf{C}[i,j] = W(j,i) - W(i,j)$.

A transition t_j is *enabled* in a marking $\mathbf{m}$ iff $\mathbf{m} \geq \mathbf{i}(t_j)$. Being enabled, t_j may *occur* (or *fire*) yielding a new marking $\mathbf{m'} = \mathbf{m} + \mathbf{C}[.,j]$ ($\mathbf{C}[.,j]$ is the jth column of $\mathbf{C}$), and this is denoted by $\mathbf{m} \xrightarrow{t_j} \mathbf{m'}$. The set of all the markings reachable from $\mathbf{m_0}$ is called *reachability set*, and is denoted by $\mathrm{RS}(\mathbf{m_0})$.

Semiflows are non-null natural annullers of $\mathbf{C}$. Right and left annullers are called T- and P-(semi)flows respectively. A semiflow is called *minimal* when its support (i.e., the set $\|\mathbf{s}\|$ of the non-zero components of vector $\mathbf{s}$) is not a proper superset of the support of any other, and the g.c.d. of its elements is one.

2.2 Previous Product Form solution results for stochastic Petri nets

A class of SPNs characterized by the fact that the stationary probability distribution of any net in this class can be factored into a product of terms has been introduced [11, 13]. Nets possessing this property are called *Product-Form Stochastic Petri Nets* (PF-SPNs) and are easily identified by the criteria proposed in [2, 7, 11, 13].

Let $\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_h}$ denote the minimal T-semiflows found from the incidence matrix. The following definitions are essential to the analysis of the SPNs that have Product Form Solution.

Definition 1. *A subset of transitions $\mathcal{T}'$ ($\mathcal{T}' \subseteq \mathcal{T}$) is said to be* closed *if $\bigcup_{t_j \in \mathcal{T}'} \mathbf{i}(t_j) = \bigcup_{t_j \in \mathcal{T}'} \mathbf{o}(t_j)$. An alternative definition of a closed set of transitions is the following: let $\mathcal{R}(\mathcal{T}') = \bigcup_{t_j \in \mathcal{T}'} \{\mathbf{i}(t_j) \cup \mathbf{o}(t_j)\}$ be the set of input and output bags for transitions in $\mathcal{T}'$. The subset of transitions $\mathcal{T}'$ is said to be* closed *if for any $\mathbf{l} \in \mathcal{R}(\mathcal{T}')$ there exists $t_i, t_j \in \mathcal{T}'$ such that $\mathbf{l} = \mathbf{i}(t_i)$ and $\mathbf{l} = \mathbf{o}(t_j)$; that is, each output bag is also an input bag for some transition in $\mathcal{T}'$, and vice-versa each input bag is also an output bag.*

Definition 2. *$\mathcal{N}$ is a Π-*net *if $\forall\, t_j \in \mathcal{T}$ there exists a minimal T-semiflow $\mathbf{x}$ such that $t_j \in ||\mathbf{x}||$, and $||\mathbf{x}||$ is a closed set.*

In other words, $\mathcal{N}$ is a Π-net if all transitions are covered by closed support minimal T-semiflows.

Example of Π-net Figure 1(a) shows a net satisfying Definition 2. We can see that there are two minimal T-semiflows $\mathbf{x_1} = [1, 0, 1, 0]$ and $\mathbf{x_2} = [0, 1, 0, 1]$, with $||\mathbf{x_1}|| = \{t_1, t_3\}$ and $||\mathbf{x_2}|| = \{t_2, t_4\}$. We can observe that $\bigcup\limits_{t_j \in ||\mathbf{x_1}||} \mathbf{i}(t_j) = \{[1, 0, 0, 0], [0, 0, 1, 0]\} = \bigcup\limits_{t_j \in ||\mathbf{x_1}||} \mathbf{o}(t_j)$ and $\bigcup\limits_{t_j \in ||\mathbf{x_2}||} \mathbf{i}(t_j) = \{[1, 1, 0, 0], [0, 0, 0, 1]\} = \bigcup\limits_{t_j \in ||\mathbf{x_2}||} \mathbf{o}(t_j)$. Both T-semiflows have closed support set. Since any transition belongs to a closed support minimal T-semiflow, this net is a Π-net.

The definition of Π-nets was originally motivated while studying the problem of finding product form solution for SPNs [2, 7, 11, 13]. More precisely, for the SPNs having the Π property, there exists a positive solution for the traffic equations (see below). In a Π-net we denote by $\mathcal{X}_c = \{\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_l}\}$ the set of closed support minimal T-semiflows. Among the minimal closed support T-semiflows, we can identify a relation that can be used to derive the PFS. Two different minimal closed support T-semiflows $\mathbf{x}'$ and $\mathbf{x}''$ are said to be *freely related*, denoted as $(\mathbf{x}', \mathbf{x}'') \in FR$, if there exist $t_j \in ||\mathbf{x}'||$ and $t_h \in ||\mathbf{x}''||$ such that $\mathbf{i}(t_j) = \mathbf{i}(t_h)$. The relation FR* is the transitive closure of FR. It is easy to see that the relation FR* yields a partitioning of the set of minimal closed support T-semiflows. Because any t_j can belong to only one FR-class, the partition of T-semiflows leads to a partition of transitions. In the following we denote by $\mathcal{C}(t_j)$ the set of the partition to which transition t belongs.

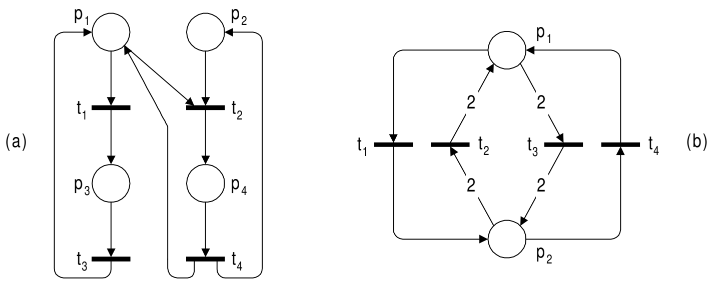

Fig. 1. Examples of Π-net s

As for Queueing Networks, PF solutions for SPN are based on the analysis of underlying Markov chains (MC). Instead of reasoning in terms of the MC with states as markings, it is more convenient to study an auxiliary MC with states being the input (or output) vectors $\mathbf{i}(t)$, called the *routing process* [11] of the SPN. The infinitesimal generator $\mathbf{Q}$ of this MC is defined by: $q(\mathbf{i}(t_j), \mathbf{o}(t_j)) = \mu(\mathbf{i}(t_j))\mathbf{P}[\mathbf{i}(t_j), \mathbf{o}(t_j)]$ with $\mu(\mathbf{i}(t_j)) = \sum_{\mathbf{i}(t_h)=\mathbf{i}(t_j)} \mu_h$. $\mathbf{P}[a, b]$ is the routing probability from $a = \mathbf{i}(t_j)$ to b: it can be computed by examining the various transitions enabled after the firing of t_j and μ_h is the usual rate of SPN transition t_h. For the sake of simplicity, we present all the results by assuming that the transition rates are marking independent. In [10] results are presented with several kinds of marking dependent transition firing rates.

The traffic equations of the routing process are the global balance equations of this MC. Denoting with $v(\mathbf{i}(t_j))$ the so-called visit-ratio to node $\mathbf{i}(t_j)$, these equations can be stated as:

$$\forall\, t_j \in \mathcal{T}, \;\; v(\mathbf{i}(t_j)) = \sum_{t_h \in \mathcal{T}} v(\mathbf{i}(t_h))\mathbf{P}[\mathbf{i}(t_h), \mathbf{i}(t_j)] \qquad (1)$$

Boucherie and Sereno [2] showed that traffic equations and structural properties of a net are closely related.

Theorem 1 (from [2]). *Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, W, \mathcal{Q}, \mathbf{m_0})$ be a SPN. There is a non null positive solution for the Traffic Equations (1) iff $\mathcal{N}$ is a Π-net.*

The existence of a positive solution for the Traffic Equations (1) is not a sufficient condition to assert a Product-Form Solution for the SPN. The following result from Coleman *et al.* [7] and [11], states that the equilibrium distribution has a product-form over the places of the SPN whenever one additional condition holds. Let us denote $f = v/\mu$ with v a solution for the traffic equations, and define the vector $\mathbf{w}_f = [w_1, \ldots, w_n]$ as

$$\mathbf{w}_f = \left[\log\left(\frac{f(\mathbf{i}(t_1))}{f(\mathbf{o}(t_1))}\right), \log\left(\frac{f(\mathbf{i}(t_2))}{f(\mathbf{o}(t_2))}\right), \cdots, \log\left(\frac{f(\mathbf{i}(t_{nt}))}{f(\mathbf{o}(t_{nt}))}\right)\right] \qquad (2)$$

There may be many functions f that derive from solutions for the traffic equations. However each one is unique up to a multiplicative constant in each FR* class. This implies that the ratio $f(\mathbf{i}(t_i))/f(\mathbf{o}(t_i))$ is invariant.

Theorem 2 (Product-Form for Equilibrium Distribution of SPN, (from [7, 11])). *Let $f = v/\mu$ with v a solution for the traffic equations. The equilibrium distribution for the SPN has the form*

$$\pi(\mathbf{m}) = \frac{1}{G}\prod_{i=1}^{np} y_i^{m_i} \qquad \forall\ \mathbf{m} \in RS(\mathbf{m_0}) \tag{3}$$

if and only if $Rank(\mathbf{C}) = Rank([\mathbf{C} \mid \mathbf{w}_f])$ where $[\mathbf{C} \mid \mathbf{w}_f]$ is the matrix $\mathbf{C}$ augmented with the row $\mathbf{w}_f$ and G a normalization constant. In this case, the np-component vector $\mathbf{r} = [\log(y_1), \ldots, \log(y_{np})]$, satisfies the matrix equation $-\mathbf{r}.\mathbf{C} = \mathbf{w}_f$.

It must be noted that, generally, the condition $Rank(\mathbf{C}) = Rank([\mathbf{C} \mid \mathbf{w}_f])$ depends on the *rates* of the transitions of the net and not only on the structure of the net.

2.3 Examples of Π-nets

Let us present two detailed examples of Π-nets. The first one complements the study of the net of Figure 1(a) and the second one shows a more complex situation about the rank condition of Theorem 2. The reader will also find an example of an unbounded Π-net in Section 4.3.

Example 1 In this example we briefly review the procedure used to obtain the equilibrium distribution for the Π-net depicted in Figure 1(a). For additional details the reader is referred to [2, 7, 11–13]. Since we know that the SPN is a Π-net, there is a solution for the Traffic equations (1):

$$\begin{array}{ll} v(\mathbf{i}(t_1)) = v(\mathbf{i}(t_3)) & \quad v(\mathbf{i}(t_3)) = v(\mathbf{i}(t_1)) \\ v(\mathbf{i}(t_2)) = v(\mathbf{i}(t_4)) & \quad v(\mathbf{i}(t_4)) = v(\mathbf{i}(t_2)) \end{array}$$

One solution is $v(\mathbf{i}(t_1)) = v(\mathbf{i}(t_3)) = v(\mathbf{i}(t_2)) = v(\mathbf{i}(t_4)) = 1$, from which we obtain $f(\mathbf{i}(t_1)) = 1/\mu_1$, $f(\mathbf{i}(t_3)) = 1/\mu_3$, $f(\mathbf{i}(t_2)) = 1/\mu_2$, and $f(\mathbf{i}(t_4)) = 1/\mu_4$. The row vector $\mathbf{w}_f$ is:

$$\begin{aligned} \mathbf{w}_f &= [\log(f(\mathbf{i}(t_1))/f(\mathbf{i}(t_3)), \log(f(\mathbf{i}(t_2))/f(\mathbf{i}(t_4))), \log(f(\mathbf{i}(t_3))/f(\mathbf{i}(t_1))), \log(f(\mathbf{i}(t_4))/f(\mathbf{i}(t_2)))] \\ &= [\log(\mu_3/\mu_1), \log(\mu_4/\mu_2), \log(\mu_1/\mu_3), \log(\mu_2/\mu_4)] \end{aligned}$$

The rank condition $Rank(\mathbf{C}) = Rank([\mathbf{C} \mid \mathbf{w}_f])$ gives us:

$$Rank\begin{pmatrix} -1 & -1 & 1 & 1 \\ 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} = Rank\begin{pmatrix} -1 & -1 & 1 & 1 \\ 0 & -1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ w_1 & w_2 & w_3 & w_4 \end{pmatrix}$$

The rank condition holds independently of the rate values because we can easily verify that $w_1 + w_3 = 0$ and $w_2 + w_4 = 0$ since $\log\left(\frac{\mu_3}{\mu_1}\right) + \log\left(\frac{\mu_1}{\mu_3}\right) = \log\left(\frac{\mu_3}{\mu_1}\frac{\mu_1}{\mu_3}\right) = \log(1) = 0$ and similarly for $w_2 + w_4 = 0$.

Since theorem 2 applies, we can obtain the expression of $\pi(\mathbf{m})$. To this end, we first solve the matrix equation $\mathbf{r}.\mathbf{C} + \mathbf{w}_f = \mathbf{0}$, that is to say:

$$
\begin{aligned}
-r_1 + r_3 + w_1 = 0 &\quad r_1 - r_3 + w_3 = 0 \\
-r_1 - r_2 + r_4 + w_2 = 0 &\quad r_1 + r_2 - r_4 + w_4 = 0.
\end{aligned}
$$

Then, setting $r_1 = r_2 = 0$, we obtain $r_3 = w_3$ and $r_4 = w_4$ from which we derive $(r_i = \log(y_i))$, $y_1 = y_2 = 1$, $y_3 = \mu_1/\mu_3$, and $y_4 = \mu_2/\mu_4$. Hence the equilibrium distribution of the SPN of Figure 1(a) is $\pi(\mathbf{m}) = \frac{1}{G}\left(\frac{\mu_1}{\mu_3}\right)^{m_3}\left(\frac{\mu_2}{\mu_4}\right)^{m_4}$.

Example 2 The SPN shown in Figure 1(b), taken form [7], represents an SPN in which the rank condition is not satisfied independently of the rate values. The incidence matrix $\mathbf{C}$ is given by $\mathbf{C} = \begin{pmatrix} -1 & 2 & -2 & 1 \\ 1 & -2 & 2 & -1 \end{pmatrix}$. This SPN is covered by four minimal T-semiflows whose support sets are $||\mathbf{x_1}|| = \{t_1, t_4\}$, $||\mathbf{x_2}|| = \{t_2, t_3\}$, $||\mathbf{x_3}|| = \{2t_1, t_2\}$, and $||\mathbf{x_4}|| = \{t_3, 2t_4\}$. Only $\mathbf{x_1}$ and $\mathbf{x_2}$ are closed, but they cover $\mathcal{T}$ so that the SPN satisfies Definition 2. Then the SPN is a Π-net and hence there exists a positive solution for the traffic equations. In particular we obtain $f(\mathbf{i}(t_i)) = \frac{1}{\mu_i}$ for $i = 1, \ldots, 4$. The vector $\mathbf{w}_f$ is given by

$$
\begin{aligned}
\mathbf{w}_f &= \left[\ \log\left(\tfrac{f(\mathbf{i}(t_1))}{f(\mathbf{i}(t_4))}\right), \log\left(\tfrac{f(\mathbf{i}(t_2))}{f(\mathbf{i}(t_3))}\right), \log\left(\tfrac{f(\mathbf{i}(t_3))}{f(\mathbf{i}(t_2))}\right), \log\left(\tfrac{f(\mathbf{i}(t_4))}{f(\mathbf{i}(t_1))}\right)\right] \\
&= \left[\log\left(\frac{\mu_4}{\mu_1}\right), \log\left(\frac{\mu_3}{\mu_2}\right), \log\left(\frac{\mu_2}{\mu_3}\right), \log\left(\frac{\mu_1}{\mu_4}\right)\right]
\end{aligned}
$$

The augmented matrix $[\mathbf{C} \mid \mathbf{w}_f]$ is row equivalent to the fully row reduced matrix $\begin{pmatrix} 1 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ w_1 & w_2 + 2w_1 & w_3 - 2w_1 & w_1 + w_4 \end{pmatrix}$.. The rank conditions are $w_2 + 2w_1 = 0$, $w_3 - 2w_1 = 0$, and $w_1 + w_4 = 0$, which implies, $\frac{f(\mathbf{i}(t_2))}{f(\mathbf{i}(t_3))}\left(\frac{f(\mathbf{i}(t_1))}{f(\mathbf{i}(t_4))}\right)^2 = 1$, $\frac{f(\mathbf{i}(t_3))}{f(\mathbf{i}(t_2))}\left(\frac{f(\mathbf{i}(t_4))}{f(\mathbf{i}(t_1))}\right)^2 = 1$, and $1 = 1$ respectively. The first and second conditions are the same and arise because there is more than one way to produce the same change of marking. Substituting for the function f, the rank condition becomes $\frac{\mu_2}{\mu_3} = \left(\frac{\mu_4}{\mu_1}\right)^2$. If this condition is met, theorem 2 applies, and, letting $y_2 = 1$ gives $y_1 = \frac{f(\mathbf{i}(t_1))}{f(\mathbf{i}(t_4))}$. Finally, $\pi_f(\mathbf{m}) = \left[\frac{\mu_4}{\mu_1}\right]^{m_1}$.

3 The class of Π-nets

In this section we are interested in structural properties of Π-nets. We present first an important result which allows one to check, in polynomial time[1] , whether a given SPN is or not a Π-net. Then, trying to position the class of Π-nets with respect to classical structural classes of PN, we show that there is no simple relation between these classes and Π-nets.

3.1 Membership problem

Algorithm Verify Π-net
$\mathcal{L} \leftarrow \mathcal{T}$
fail $\leftarrow$ *false*
repeat
  let $t \in \mathcal{L}$
  $\mathcal{A} \leftarrow \{t\}$
  $In \leftarrow \{\mathbf{i}(t)\}$
  $Out \leftarrow \{\mathbf{o}(t)\}$
  while $\exists t' \in \mathcal{L}$ s.t. $\mathbf{i}(t') \in Out$ **do**
    $\mathcal{A} \leftarrow \mathcal{A} \bigcup \{t'\}$
    $\mathcal{L} \leftarrow \mathcal{L} \setminus \{t'\}$
    $In \leftarrow In \bigcup \{\mathbf{i}(t')\}$
    $Out \leftarrow Out \bigcup \{\mathbf{o}(t')\}$
  endwhile
  fail $\leftarrow (In \neq Out)$
  /* if not*fail* then $\mathcal{A}$ is a FR^* class */
until $\mathcal{L} = \emptyset$ or *fail*
/* *fail* is true iff the net is not a Π-net */

From the definition of Π-nets we can decide if a given net falls in this class. The problem that arises is the complexity of a straightforward application of Definition 2 because the number of minimal T-semiflows can be exponential in the number of transitions (e.g., [17]). We present now an algorithm that allows to recognize whether a net is a Π-net in polynomial time. The soundness of the algorithm is based on the following lemma (see [2] for the proof).

Lemma 1. *If $\mathbf{x}$ is a closed support minimal T-semiflow then (i) for each transition $t_i \in ||\mathbf{x}||$, $\mathbf{x}[i] = 1$ ($\mathbf{x}[i]$ is the i-th component of $\mathbf{x}$). (ii) $||\mathbf{x}||$ may be ordered as $\left\{t_{j_0}, t_{j_1}, \ldots, t_{j_{h-1}}\right\}$ such that $\mathbf{o}(t_{j_i}) = \mathbf{i}(t_{j_{i+1\, mod\, h}})$ (for $i = 0, 1, \ldots, h-1$), and $l \neq l' \Rightarrow \mathbf{i}(t_{j_l}) \neq \mathbf{i}(t_{j_{l'}})$.*

Algorithm for Π-net membership The previous lemma states that a closed support minimal T-semiflow can be seen as a cycle of transitions $t_{j_0}, t_{j_1}, \ldots, t_{j_{h-1}}$ such that $\mathbf{o}(t_{j_i}) = \mathbf{i}(t_{j_{i+1\, mod\, h}})$ (for $i = 0, 1, \ldots, h-1$). The algorithm **Verify Π-net** exploits this feature for checking if a net is a Π-net.

[1] Unless explicitly mentioned, all complexity results in the paper are w.r.t. the size of the net, i.e. the number of places, transitions, arcs and the binary representation of valuations.

We point out that the algorithm yields a covering set of closed support minimal T-semiflows (if the SPN is a Π-net). From then we can derive the routing probabilities and the partitions of the set of transitions $\mathcal{T}$ into FR*-classes.

From simple considerations we can see that both the inner and the outer cycles require $O(|\mathcal{T}|)$ steps. Hence the complexity of the algorithm that allows to recognize if a given net satisfy Definition 2 requires $O(|\mathcal{T}|^2)$ steps.

3.2 *Π*-nets and other classes of PN

As usual for Petri net models, it is interesting to examine whether it is possible to structurally characterize behavioural properties of these nets and to deduce efficient checking of these properties. Since this is the case for some well known subclasses of nets, we first recall such subclasses and we compare Π-nets with them. For completeness, results include the class of $\overline{\Pi}$-nets which are introduced in section 3.3.

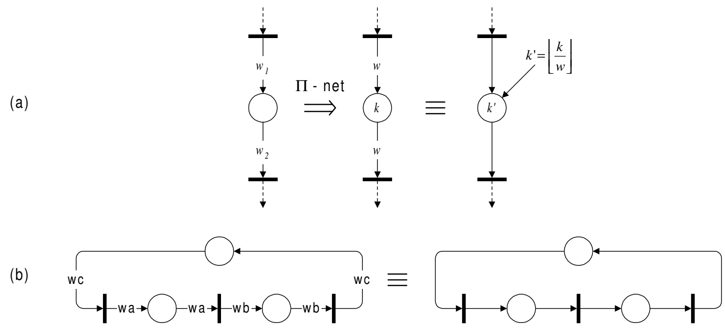

Fig. 2. Conversion of WTS Π-nets into MGs

The following classes of Petri nets are particularly interesting for the analysis of behavioural properties:

- A state machine (SM) is a Petri net with binary valuations where any transition has exactly one input and one output place.
- A marked graph (MG) is a Petri net with binary valuations where any place has exactly one input and one output transition.
- A weighted transition system (WTS) is a Petri net where any place has exactly one input and one output transition (MG are special case of WTS).
- An extended free-choice net (EFC) is a Petri net with binary valuations where two transitions, sharing an input place, have the same set of input places.

Proposition 1. *Comparing Π-nets with some classical subclasses of Petri nets, we have:*

- *If $\mathcal{N}$ is a WTS and a Π-net, then it is behaviourally equivalent to a MG.*

- *Every SM is a Π-net (and even a $\overline{\Pi}$-net).*
- *There are MGs which are not Π-nets.*
- *There are Π-nets (and even $\overline{\Pi}$-nets) which are non EFC nets.*

Proof. Figure 2 explains the conversion from a WTS Π-net to a MG: in (a) we change the weights of arcs connecting isolated places ($k = w_1 - w_2$); in (b), we observe that any weighted Π-cycle is just equivalent to an ordinary cycle.
As a straightforward consequence of the definitions, every SM is a Π-net. In any SM, $\mathbf{r}$ vectors are $\mathbf{1_p}$: null components except on component p, input or output place of a transition t. Taking $\mathbf{a_r} = \mathbf{r}$ for each $\mathbf{r}$, we see that a SM is also a $\overline{\Pi}$-net (see below for definition of $\overline{\Pi}$-nets).
A net with an idle place followed by a parbegin-parend with intermediate action is a MG but not a Π-net. Note however, that any Π-net MG is a union of disjoint cycles, hence a $\overline{\Pi}$-net.
Finally, we will see that the net of Example 1 (Figure 1(a)) is a $\overline{\Pi}$-net, and it is clearly not an EFC.

3.3 $\overline{\Pi}$-nets

In this section we define the class of $\overline{\Pi}$-nets which are exactly the set of Π-nets having a PF solution *for any stochastic specification* in contrast with previous results whose criteria are dependent on rates of transitions (see Example 2). Moreover, we introduce a more general dependency of the firing rates of transitions with respect to the global marking of the net system.

Definition of $\overline{\Pi}$-nets Criteria found by several authors since the late 80's for PF solution of SPN are only sufficient conditions, and moreover, they are made up of structural conditions and conditions on the stochastic parameters of the SPN. In search of a pure structural characterization of PF solution SPN, we were led to fully reconsider the concept of "virtual client state" of a Π-net system in the context of routing processes and to deeply analyze how to characterize these states. In previous works, T-semiflows identify concurrent "virtual clients" activities. These activities are "synchronized" by conflicting resources allocation, that is shared input places of transitions. For what concerns places, they are usually interpreted either as specific resources or as clients. But, indeed, this interpretation does not allows us to express the PF property at a structural level because virtual client states do not reduce to place markings, even in a Π-net. For instance, in the example net 1 (figure 1(a)), we may think of t_1, p_3, t_3 as batch jobs processing (activity 1), and of t_2, p_4, t_4, p_2 as interactive work of users (activity 2). The place p_1, modelling processor resources, cannot, alone, characterize the "idle" batch jobs state. This is the crucial point: in a Π-net, we have no information about the state of the virtual clients in the net system and this is the main reason which prevents us to state a necessary and sufficient condition for the existence of a PF solution. Actually, we have found that virtual client states are characterized by a relation $\mathbf{v}.\mathbf{C} = \mathbf{r}$, where $\mathbf{v}$ is a vector on places and $\mathbf{r}$ is a vector such that $\mathbf{r}[t] = 1$ if t adds a client to the "state", $\mathbf{r}[t] = -1$ if t

removes a client and $\mathbf{r}[t] = 0$ otherwise. The $\overline{\Pi}$-net property expresses, by means of rational vectors $\mathbf{a_r}$, the relation which must hold between virtual clients states of a Π-net and input/output vectors of the net, to ensure that this Π-net has a PF solution,

Moreover, this explicit relation on states of virtual clients allows us to model the dependency of the firing rate of a transition t_j with respect to the global state of the system in parts (activities) of the net not related to the input/output vectors of transitions belonging to $\mathcal{C}(t_j)$. This kind of dependency, introduced by functions $\rho_{\mathcal{C}(t_j)}$ in the definition below, cannot be taken into account in the framework of Π-nets.

For the rest of this section, we set: ${}^\bullet\mathbf{r} = \{t_j \in \mathcal{T} \mid \mathbf{o}(t_j) = \mathbf{r}\}$ and $\mathbf{r}^\bullet = \{t_j \in \mathcal{T} \mid \mathbf{i}(t_j) = \mathbf{r}\}$ for every $\mathbf{r} \in \mathcal{R}(\mathcal{T})$.

Definition 3 ($\overline{\Pi}$-net). *A $\overline{\Pi}$-net (*restricted Π-net*) is a Π-net such that for every $\mathbf{r} \in \mathcal{R}(\mathcal{T})$, there exists $\mathbf{a_r} \in \mathbb{Q}^{|\mathcal{P}|}$ such that*

$$\mathbf{a_r}.\mathbf{C}[\mathcal{P}, j] = \begin{cases} 1 & \textit{if } t_j \in {}^\bullet\mathbf{r} \\ -1 & \textit{if } t_j \in \mathbf{r}^\bullet \\ 0 & \textit{otherwise} \end{cases}$$

where $\mathbf{C}$ is the incidence matrix of the net (note that this excludes transitions t_h with $\mathbf{i}(t_h) = \mathbf{o}(t_h)$).
The firing rate of a transition t_j of a $\overline{\Pi}$-net system in the marking $\mathbf{m}$ is given by

$$\mu(t_j, \mathbf{m}) = \mu(\mathbf{i}(t_j)).\rho_{\mathcal{C}(t_j)}\left((\mathbf{a_{r''}}.\mathbf{m})_{\mathbf{r''}\notin\mathcal{C}(t_j)}\right).\mathbf{P}[\mathbf{i}(t_j), \mathbf{o}(t_j)] \quad (4)$$

Positive, real valued functions $\rho_{\mathcal{C}(t_j)}\left((\mathbf{a_{r''}}.\mathbf{m})_{\mathbf{r''}\notin\mathcal{C}(t_j)}\right)$ make possible a homogeneous dependency of the transitions of the component $\mathcal{C}(t_j)$ w.r.t. the state of the virtual clients in the other components, given by the $\mathbf{a_{r''}}.\mathbf{m}$ (see example below).

Note that the computation of the rational vectors $\mathbf{a_r}$ (or else the proof that there are no such $\mathbf{a_r}$), may be achieved in polynomial time with respect to the size of the net through a usual Gaussian elimination (but restricted to rational numbers).

The net of Example 2 is an example of a Π-net which is not a $\overline{\Pi}$-net(see its incidence matrix in Section 2.3). Let us set $\mathbf{r_1} = \{p_1\}$, so that ${}^\bullet\mathbf{r_1} = \{t_4\}$ and $\mathbf{r_1^\bullet} = \{t_1\}$. If we try to define the vector $\mathbf{a_{r_1}} = [a, b]$, we get $a - b = 1$ (since $t_4 \in {}^\bullet\mathbf{r_1}$) and $a - b = 0$ (since $t_2 \notin {}^\bullet\mathbf{r_1} \bigcup \mathbf{r_1^\bullet}$). Hence, $\mathbf{a_{r_1}}$ does not exist and this SPN is not a $\overline{\Pi}$-net. In fact, t_1 and t_3 have proportional input and output bags but belong to different T-semiflows and no distinction between these transitions is possible from $\mathbf{r_1} = [1, 0, 0, 0]$.

The Π-net of Example 1 (see Section 2.3) is a $\overline{\Pi}$-net. We have four input vectors $\mathbf{r}$, belonging to two classes: $\mathcal{C}_1 = \{\mathbf{r_1} = [1, 0, 0, 0], \mathbf{r_3} = [0, 0, 1, 0]\}$, $\mathcal{C}_2 = \{\mathbf{r_2} = [1, 1, 0, 0], \mathbf{r_4} = [0, 0, 0, 1]\}$. The $\mathbf{a_r}$ vectors are

$$\mathbf{a_{r_1}} = [0, 0, -1, 0] \quad \mathbf{a_{r_3}} = [0, 0, 1, 0]$$
$$\mathbf{a_{r_2}} = [0, 0, 0, -1] \quad \mathbf{a_{r_4}} = [0, 0, 0, 1]$$

Let us assume that the rate of t_3 depends on the load of t_4 in such a way that if the marking of p_4 is greater than K_4, t_3 cannot fire (because no more resource is available for instance). Moreover, suppose that the rate of t_4 decreases linearly from μ_M to μ_m with the marking of p_4 varying from 0 to K_4. Then we can define

$$\rho_{\mathcal{C}(t_3)}\left((\mathbf{a}_{\mathbf{r}''}.\mathbf{m})_{\mathbf{r}''\notin\mathcal{C}(t)}\right) = \begin{cases} 0 & \text{if } \mathbf{m}[p_4] \geq K_4 \\ \frac{\mu_m-\mu_M}{K_4}.\mathbf{m}[p_4] + \mu_M & \text{if } 0 \leq \mathbf{m}[p_4] < K_4 \end{cases}$$

since $\mathbf{a}_{\mathbf{r_4}}.\mathbf{m} = \mathbf{m}[p_4]$ and we have still a PF steady state distribution.

Due to lack of space, we present in the rest of the paper, results without functions $\rho_{\mathcal{C}(t)}$. The reader will find full version of the results in the technical report [10].

Sufficient condition for PF-SPN We first establish a sufficient condition for a Π-net to have a PF steady state distribution, *whatever* the parameters (i.e. rates of transitions) of the stochastic specification of the SPN.

Theorem 3. *Let* $(\mathcal{P},\mathcal{T},W,\mathcal{Q},\mathbf{m_0})$ *be a* $\overline{\Pi}$*-net. Then,* for any transition rates, *the steady state distribution of the SPN has the product form*

$$\pi(\mathbf{m}) = \frac{1}{G}.\prod_{\mathbf{a_r}\in\mathcal{R}(\mathcal{T})}\left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})}\right)^{\mathbf{a_r}.\mathbf{m}} \qquad \forall\ \mathbf{m}\in RS(\mathbf{m_0}), \tag{5}$$

where G is a normalization constant and v is a solution of Equations (1).

Let us remark that this product form expression induces, of course, a product form with respect to $\mathbf{m}$, since:

$$\prod_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}\left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})}\right)^{\mathbf{a_r}.\mathbf{m}} = \prod_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}\prod_{p_i\in\mathcal{P}}\left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})}\right)^{\mathbf{a_r}[i].\mathbf{m}[i]} = \prod_{p_i\in\mathcal{P}}\left(\prod_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}\left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})}\right)^{\mathbf{a_r}[i]}\right)^{\mathbf{m}[i]}$$

Sketch of proof We give only a sketch of the proof (see [10] for a detailed proof). The starting point is the so-called the Group Local Balance Equation for a marking $\mathbf{m}$ with respect to a given vector $\mathbf{r}$ which is a splitting of the equilibrium (Chapman-Kolmogorov) equations of the Markov chain with markings as states:

$$\pi(\mathbf{m})\sum_{t_j\in\,\mathbf{r}^\bullet} q(\mathbf{m},\mathbf{m}-\mathbf{i}(t_j)+\mathbf{o}(t_j)) = \sum_{t_h\in{}^\bullet\mathbf{r}}\pi(\mathbf{m}+\mathbf{i}(t_h)-\mathbf{o}(t_h))q(\mathbf{m}+\mathbf{i}(t_h)-\mathbf{o}(t_h),\mathbf{m}] \tag{6}$$

Then using the expression of the rates q, we introduce the proposed expression and after simplification, we get:

$$\mu(\mathbf{r}) = \sum_{t_h\in{}^\bullet\mathbf{r}}\prod_{\mathbf{r}'\in\mathcal{R}(\mathcal{T})}\left(\frac{v(\mathbf{r}')}{\mu(\mathbf{r}')}\right)^{\mathbf{a}_{\mathbf{r}'}.(\mathbf{i}(t_h)-\mathbf{o}(t_h))}\mu(\mathbf{i}(t_h))\mathbf{P}[\mathbf{i}(t_h),\mathbf{r}]. \tag{7}$$

From $\mathbf{i}(t_h)-\mathbf{o}(t_h) = -\mathbf{C}[\mathcal{P},h]$ and the definition of $\mathbf{a_r}$, (7) can be shown equivalent to the Traffic Equations (1). ∎

Necessary condition for PF-SPN The result of this section proves that the concept of $\overline{\Pi}$-net is the adapted one to capture the existence of a product form like the one of Theorem 3 *for any stochastic specification* of a Π-net. Combining Theorems 3 and 4, the "$\overline{\Pi}$-net property" appears as a necessary and sufficient structural condition for a net to have a product form steady state distribution *for any transition rates.*

Theorem 4. *Let* $(\mathcal{P}, \mathcal{T}, W, \mathcal{Q}, \mathbf{m_0})$ *be a* Π*-net and* v *a solution of the Traffic Equations. If there is a family* $(\mathbf{a_r})_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}$ *of rational vectors such that the distribution*

$$\pi(\mathbf{m}) = \frac{1}{G} \cdot \prod_{\mathbf{r}\in\mathcal{R}(\mathcal{T})} \left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})}\right)^{\mathbf{a_r}.\mathbf{m}} \qquad \forall\ \mathbf{m} \in RS(\mathbf{m_0}),$$

satisfies the Group Local Balance Equations (6) for any $(\mu(\mathbf{r}))_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}$*, then we have*

$$\mathbf{a_r} \cdot \mathbf{C}[\mathcal{P}, j] = \begin{cases} \ \ 1 & \textit{if } t_j \in {}^\bullet\mathbf{r} \\ -1 & \textit{if } t_j \in \mathbf{r}^\bullet \\ \ \ 0 & \textit{otherwise} \end{cases}$$

Sketch of proof (see [10] for a detailed proof). The Group Local Balance Equations for a given $\mathbf{m}$ with respect to a given $\mathbf{r}$ are (see (7))

$$\mu(\mathbf{r}) = \sum_{t_j\in{}^\bullet\mathbf{r}} \prod_{\mathbf{r'}\in\mathcal{R}(\mathcal{T})} \left[\frac{v(\mathbf{r'})}{\mu(\mathbf{r'})}\right]^{-\mathbf{a_{r'}}.\mathbf{C}[\mathcal{P},j]} \mu(\mathbf{i}(t_j))\mathbf{P}[\mathbf{i}(t_j), \mathbf{r}] \qquad (8)$$

since $a_{\mathbf{r'}}.(\mathbf{i}(t_j) - \mathbf{o}(t_j)) = -\mathbf{a_{r'}}.\mathbf{C}[\mathcal{P}, j]$.
The idea is to express (8) as a multi-variables identically null "polynom" (i.e. extension of multi-variables polynom, with real valued exponents instead of integer) on $\mathbb{R}^+$ and to deduce the claimed properties of the $\mathbf{r}$ vectors from properties of the coefficients of this "polynom". To this end, we introduce the vectors with np components $\gamma(t_j)$ and $\gamma_\mathbf{0}$ in the following way:

$$\gamma(t_j)[\mathbf{r'}] = \begin{cases} a_{\mathbf{r'}}.\mathbf{C}[\mathcal{P}, j] & \text{if } \mathbf{r'} \neq \mathbf{i}(t_j) \\ a_{\mathbf{r'}}.\mathbf{C}[\mathcal{P}, j] + 1 & \text{if } \mathbf{r'} = \mathbf{i}(t_j) \end{cases} \qquad \text{and} \qquad \gamma_\mathbf{0}[\mathbf{r'}] = \begin{cases} 1 & \text{if } \mathbf{r'} = \mathbf{r} \\ 0 & \text{otherwise} \end{cases}$$

Using these vectors, transformation of Equation (8) provides a "polynom" with variables $\mu(\mathbf{r'})$. Via a technical result, it can then be shown that for all t_j, the set $\{t_j \in {}^\bullet\mathbf{r} \mid \gamma(t_j) = \gamma \neq \gamma_\mathbf{0}\}$ is empty, so that $\forall t_j \in {}^\bullet\mathbf{r},\ \gamma(t_j) = \gamma_\mathbf{0}$. The result then follows from the evaluation of the numbers $a_{\mathbf{r'}}.\mathbf{C}[\mathcal{P}, j]$. ∎

4 Functional properties of PF-SPN

Although Π-nets and $\overline{\Pi}$-nets are not easily comparable to standard classes of PN, they nevertheless enjoy specific qualitative properties. This section first reviews liveness and deadlock freeness in Π-nets; second, some results about the complexity of the reachability and liveness in Π-nets and $\overline{\Pi}$-nets are presented. Finally, we expose results about the characterization of reachable markings in Π-nets. Since we need to distinguish between structural and behavioural properties of (S)PN, in this section, we denote by $\mathcal{N} = (\mathcal{P}, \mathcal{T}, W, \mathcal{Q})$ a SPN and by $\Sigma = (\mathcal{N}, \mathbf{m_0})$ a marked SPN (also called SPN *system*) with initial marking $\mathbf{m_0}$.

4.1 Some behavioural properties of Π-nets

Liveness is an important property of Petri net systems. Due to importance of T-semiflows in Π-nets, it is not surprising that liveness in Π-nets systems enjoys particular properties that we present below together with related results. The following lemma is a direct consequence of Proposition 1.

Lemma 2. *Let $\Sigma = (\mathcal{N}, \mathbf{m_0})$ be a Π-system. If $t \in \mathcal{T}$ is enabled at $\mathbf{m} \in \mathrm{RS}(\mathbf{m_0})$, then,*
(1) all transitions of all minimal closed support T-semiflows to which t belongs can be fired.
(2) there is a firing sequence that fires all the remaining transitions in the FR*-class of t.

Proposition 2. *Let $\Sigma = (\mathcal{N}, \mathbf{m_0})$ be a Π-system.*

1. *If $\exists\, t \in \mathcal{T}$, enabled at $\mathbf{m_0}$ then Σ is deadlock-free (DF).*
2. *Σ is reversible.*
3. *$\mathcal{N}$ is structurally live (SL).*
4. *If there is an enabled transition in any FR*-class in the initial marking, then Σ is live. The converse is false.*
5. *If Σ is live then $\Sigma' = (\mathcal{N}, \mathbf{m}'_0)$ with $\mathbf{m_0} \leq \mathbf{m}'_0$ is live too (i.e., liveness is monotonic w.r.t the initial marking in the net).*

Proof. We only give the detailed proof of (2).
If $\mathbf{m_0}$ is not a deadlock marking, for any $\mathbf{m} \in \mathrm{RS}(\mathbf{m_0})$ there is a finite firing sequence $\sigma = t_{\delta_1}, t_{\delta_2}, \ldots, t_{\delta_l}$ such that $\mathbf{m_0}[t_{\delta_1}\rangle\mathbf{m_1} \ldots \mathbf{m_{l-1}}[t_{\delta_l}\rangle\mathbf{m}$. Now we prove that there is a finite firing sequence η such that $\mathbf{m}[\eta\rangle\mathbf{m_0}$. Let $\mathbf{x}$ be a closed support T-semiflow (not necessarily minimal) such that $\mathbf{x} \geq \boldsymbol{\sigma}$. Since $\mathbf{x}$ is a linear combination of minimal closed support T-semiflows, it follows from Lemma 2 that from $\mathbf{m}$, $\mathbf{x} - \boldsymbol{\sigma}$ must be firable and hence $\mathbf{m}[\mathbf{x} - \boldsymbol{\sigma}\rangle\mathbf{m_0}$.

Reverse of Π-net Finally, the next proposition addresses properties of the reverse net of Π-nets. The reverse net of a Petri net $\mathcal{N} = (\mathcal{P}, \mathcal{T}, W)$ is $\mathcal{N}^{(-1)} = (\mathcal{P}, \mathcal{T}, W^{(-1)})$, that is, the net with same places and transitions, but reversed arcs ($W^{(-1)}(i,j) = W(j,i)$). Note that $(\mathcal{N}^{(-1)})^{(-1)} = \mathcal{N}$ and that the incidence matrix of $\mathcal{N}^{(-1)}$ is $-C$.

Proposition 3. *Let $\mathcal{N}$ be a Π-net, $\Sigma = (\mathcal{N}, \mathbf{m_0})$, and $\Sigma^{(-1)} = (\mathcal{N}^{(-1)}, \mathbf{m_0})$.*

1. *The reverse of a Π-net (resp. $\overline{\Pi}$-net) is a Π-net (resp. $\overline{\Pi}$-net).*
2. *Σ is deadlock free iff $\Sigma^{(-1)}$ is deadlock free.*
3. *The reachability graph of $\Sigma^{(-1)}$ is the reverse of the reachability graph of Σ.*

4. Σ is live iff $\Sigma^{(-1)}$ is live.

Proof. For space savings, we only develop proof of (3). If $\mathbf{m_0}$ is not a deadlock marking, then from Proposition 2 (1) and (2), Σ is reversible. But in any reversible Petri net, the announced property holds. Indeed, we have first $\mathrm{RS}(\Sigma^{(-1)}) = \mathrm{RS}(\Sigma)$. Let $\mathbf{m} \in \mathrm{RS}(\Sigma)$. Since Σ is reversible, there is a firing sequence τ such that $\mathbf{m}[\tau\rangle\mathbf{m_0}$. Therefore, $\mathbf{m_0}[\tau^{(-1)}\rangle\mathbf{m}$ in $\Sigma^{(-1)}$ where $\tau^{(-1)}$ is τ with "reversed" transitions. Now, let $\mathbf{m}[t\rangle\mathbf{m'}$ in Σ. We have $\mathbf{m'} \in \mathrm{RS}(\Sigma^{(-1)})$ and, obviously, $\mathbf{m'}[t^{(-1)}\rangle\mathbf{m}$. We have proven that the reverse of the reachability graph Σ is a partial graph of $\Sigma^{(-1)}$. The result follows, applying the same proof to $\Sigma^{(-1)}$.

4.2 Complexity of liveness and reachability problems for Π-nets and $\overline{\Pi}$-nets

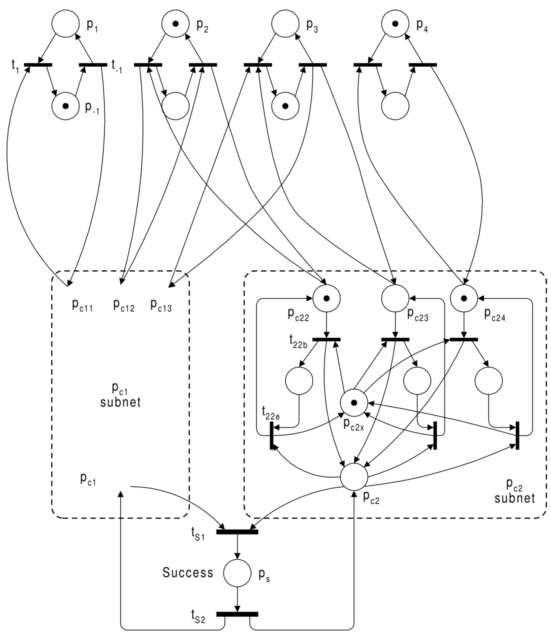

Fig. 3. Reduction of 3SAT to liveness in 1-safe $\overline{\Pi}$-nets

Condition (4) in proposition 2 is only a sufficient condition. In fact, checking liveness seems no more easy for Π-nets, and even 1-safe[2] $\overline{\Pi}$-nets, than for many

[2] A 1-safe marked Petri net is a (bounded) marked net with at most one token in every place of every reachable marking

other classes of Petri nets. We have shown in Section 3.1 that the complexity of the computation of FR*-Classes is polynomial time. But checking liveness requires to verify that each FR*-class is live. If some FR*-class is not *initially* firable, this is still a very complex problem. Indeed the next lemma gives some insight into this point. We recall that for general Petri nets, Lipton's result [16] implies a $2^{O(\sqrt{n})}$ lower bound space complexity for the liveness problem (see [9, 8] for recent surveys on decidability problems for Petri nets). In fact, we are able to give more precise results, although the exact complexity of the reachability/liveness for $\overline{\Pi}$-nets still remains an open problem.

It has been shown in [6] that the liveness problem for 1-safe nets is PSPACE complete. The next lemma gives a lower bound of the problem for 1-safe Π-nets.

Proposition 4. *The liveness problem for 1-safe $\overline{\Pi}$-nets is NP-hard.*

Proof. To prove it, we reduce in polynomial time the 3SAT problem to the liveness problem for $\overline{\Pi}$-nets, following the idea first presented in [14]. The 3SAT problem is a well known NP-complete problem. We have K logical formulae $C_1, \cdots, C_K$, each one being a disjunction of three boolean variables v_i or their negation $(-v_i)$, from a set of I variables: for instance, $C_k = v_1 \vee -v_3 \vee v_6$. The 3SAT problem is: is there a set of values for $v_1, \cdots, v_I$ such that $C_1 \wedge C_2 \wedge \cdots \wedge C_K$ is true? We explain the reduction through the example $C_1 = v_1 \wedge -v_2 \wedge v_3$, $C_2 = v_2 \wedge v_3 \wedge v_4$ $(K = 2, I = 4)$ (Figure 3).

For each variable v_i, we have two places p_i and p_{-i} and two transitions t_i and t_{-i}. Arcs between places and transitions for v_i are as indicated in the figure. We have also K sets of places $p_{Ck,i}$ (the introduction of several places for each C formulae ensures 1-safeness). If v_i is in C_k (like v_2 and C_2) there is an arc from t_{-i} to $p_{Ck,i}$ and one arc from $p_{C_k i}$ to t_i. In contrast if $-v_i$ is in C_k (like $-v_2$ in C_1), these arcs are reversed. Otherwise, there is no arc between t_i, t_{-i} and place $p_{C_k i}$. Places detailed in the right dotted part ensure that the place p_{Ck} will contain at most one token (p_{C2x} is a mutex place). Finally, we have one transition t_{s1} (for Success) and we added place p_s and transition t_{s2} to have a $\overline{\Pi}$-net and not only a Π-net.

We can easily verify that the net is a 1-safe $\overline{\Pi}$-net. The initial marking is chosen as follows: if v_i is true, there is one token in p_i and one token in place $p_{C_k i}$ if v_i is in C_k; if v_i is false, there is one token in p_{-i} and one token in place $p_{C_k i}$ if $-v_i$ is in C_k. In our example, we take $v_1 = v_3$ =false, $v_2 = v_4$ =true. Clearly the formula is true for a given set of boolean values of variables if the transition t_{s1} is live and the same for the reachability of a marking with one token in p_s.

Thus, there is still an open problem for $\overline{\Pi}$-nets since the upper bound of complexity for general Petri nets is in PSPACE. By contrast, the next proposition provides an exact characterization of the complexity of the problems for (1-safe) Π-nets. This distinctive result strengthens the specific character of the Π-nets class.

Proposition 5. *(1) The liveness and the reachability problems for 1-safe Π-nets are PSPACE complete.*
(2) The reachability problem for Π-nets is EXPSPACE-complete.

Proof. Due to lack of space, we only address the claim (2). For symmetric nets systems, we know [5, 18] that the reachability problem is EXPSPACE complete. A net is symmetric iff for every transition t, there is a "reverse" transition t' whose firing "undoes" the effect of the firing of t, i.e., the input places of t are the output places of t' and vice versa. Symmetric nets are clearly Π-nets. Thus, the reachability problem is EXPSPACE-hard for Π-nets. But any Π-net defines implicitly a symmetric net: for any transition t, we may add a reverse transition t' without changing the resulting reachability graph, because the closed T-semiflow (without t) of transitions to which t belongs acts exactly as t' when fired in a cyclic way. Thus, the reachability problem for Π-nets is reducible to the one for symmetric nets, hence in EXPSPACE, and finally EXPSPACE-complete.

4.3 Algebraic properties of PF-SPN

The availability of a product form equilibrium distribution allows the development of computational algorithms that are analogous to those developed for product form solution queueing networks (e.g, [3, 4, 21]). For instance proposals for algorithms for the computation of performance measures throughout the normalization constant calculus can be found in [7, 23]. In [1] a set of *Arrival Theorems*, similar to the analogous results developed for product form solution queueing networks [25] was proven, leading to a Mean Value Analysis (MVA) for the computation of performance measures for PF-SPNs. MVA for SPNs was also studied in [24].

This last section discusses reachability markings properties related to the solution of PF-SPN. For the development of computational algorithms for PF-SPN, the reachability set (RS) of the SPN must be partitioned according to certain criteria depending on the particular algorithm. For instance, the normalization constant computation algorithm requires a partitioning of the reachability set that groups together all the markings with a constant number of tokens in a given place. It is then important to know if reachable markings of a Π-net system may be characterized, among all markings, by some specific criterion based on their value and structural elements of the net. The most common such criteria are the so-called *state equation* and the one *based on the minimal P-semiflows* of the net. The difficulty then lies in the quality of those criteria, i.e. whether they allow to select all reachable markings and, *only* reachable markings.

Let us recall that the *state equation* $\mathbf{m} = \mathbf{m_0} + \mathbf{C}.\boldsymbol{\sigma}$ is an algebraic equation that gives a necessary condition for a marking to be reachable. The set of vectors $\mathbf{m} \in \mathbb{N}^{np}$ such that $\exists\ \boldsymbol{\sigma} \in \mathbb{N}^{nt} : \mathbf{m} = \mathbf{m_0} + \mathbf{C}.\boldsymbol{\sigma}$ is called the Potential Reachability Set (PRS) of the net. Obviously, $\mathrm{RS}(\mathbf{m_0}) \subseteq \mathrm{PRS}(\mathbf{m_0})$. In the literature, there are several proposals of computational algorithms for PF-SPN. They use a reachability characterization based on the minimal P-semiflows. Therefore, another set of "potential" markings has been defined. Let $\mathbf{B}$ be the matrix whose rows are the set of minimal P-semiflows of the net. The *Potential Reachability Set with respect to* $\mathbf{B}$ is the set $\mathrm{PRS}^{\mathbf{B}}(\mathbf{m_0}) = \{\mathbf{m} \mid \mathbf{B}.\mathbf{m} = \mathbf{B}.\mathbf{m_0}\}$. Clearly, $\mathrm{PRS}(\mathbf{m_0}) \subseteq \mathrm{PRS}^{\mathbf{B}}(\mathbf{m_0})$ since $\mathbf{B}.\mathbf{C} = 0$.

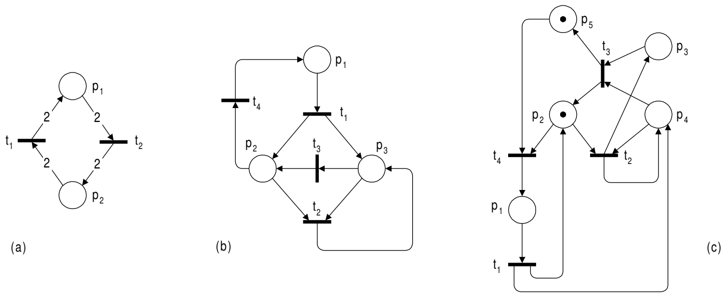

Fig. 4. Π-net and potential reachability: (a) $\mathrm{PRS}(\mathbf{m_0}) \neq \mathrm{PRS}^{\mathbf{B}}(\mathbf{m_0})$, unbounded (b) and bounded (c) Π-systems with spurious marking

An unreachable marking belonging to one of these PRS is called a *spurious marking* (see [27] for a detailed study of several kinds of PRS). We show below that, unfortunately, none of these two characterizations is able to capture all the peculiarities of PF-SPN, that is to say that there are Π-net with spurious markings for PRS (thus for $\mathrm{PRS}^{\mathbf{B}}$).

First we may have $\mathrm{PRS}(\mathbf{m_0}) \neq \mathrm{PRS}^{\mathbf{B}}(\mathbf{m_0})$ in Π-systems. This happens even in such simple case as the Π-cycle of Figure 4(a): the dead marking $\mathbf{m_1} = [1,1]^T$ has the same dot-product with the P-semiflow $\mathbf{Y} = [1,1]^T$ as the live one $\mathbf{m_0} = [2,0]$ although there is no $\boldsymbol{\sigma} \in \mathbb{N}^{nt}$ satisfying the state equation $\mathbf{m_1} = \mathbf{m_0} + \mathbf{C}.\boldsymbol{\sigma}$.

For what concerns the characterization of the reachability set of a Π-net in terms of potential reachability set, the proposition below (we omit the proof for sake of place) provides a rather positive result, but we give next, two examples which prove that properties of Π-nets are not strong enough to prevent the existence of spurious markings.

Proposition 6. *With respect to the state equation,*
(1) The potential reachability graph of $(\mathcal{N}, \mathbf{m_0})$ *is equal to the reverse of the potential reachability graph of* $(\mathcal{N}^{-1}, \mathbf{m_0})$.
(2) Spurious markings (if they exist) cannot be transient, i.e., if $\mathbf{m} \in \mathrm{PRS}(\mathbf{m_0}) \setminus \mathrm{RS}(\mathbf{m_0})$*, then there is no firing sequence* $\boldsymbol{\sigma}$ *such that* $\mathbf{m}[\sigma\rangle\mathbf{m'}$ *with* $\mathbf{m'} \in \mathrm{RS}(\mathbf{m_0})$.

The net of Figure 4(b) gives the first negative result. For the unbounded Π-net it is possible to see that $\mathbf{m} = [0,0,0]$ is a spurious marking. We can see that for any initial marking $\mathbf{m_0} = [k_1, k_2, k_3]$, $\mathbf{m_0}[t_1^{k_1}\rangle\mathbf{m_1} = [0, k_1+k_2, k_1+k_3][t_2^{k_1+k_2}\rangle\mathbf{m_2} = [0,0,2k_1+k_2+k_3]$. Setting $k = 2k_1+k_2+k_3$ we have $\mathbf{m_2}[(t_3t_2)^{k-1}\rangle\mathbf{m_3} = [0,0,1]$. Now "firing" t_2t_3 the null marking is spuriously reached.

The net of Figure 4(c) gives another and definitive negative result. This Π-net is bounded but it is possible to see that $\mathbf{m} = [0,0,1,0,1]$ is a spurious marking. Indeed, from the initial marking $\mathbf{m_0} = [0,1,0,0,1]$ and with the "firing" of t_2 we obtain the marking $[0,0,1,0,1]$ that it is not a reachable marking. Hence we have $\mathbf{m} \in \mathrm{PRS}(\mathbf{m_0})$ but $\mathbf{m} \notin \mathrm{RS}(\mathbf{m_0})$.

5 Conclusion

SPN with PF solution have been introduced some years ago as an extension of closed form solution methods of QN to SPN which allow to model systems with more complex synchronization schemes. In this paper we have presented four groups of new results giving a better insight in PF-SPN and allowing an efficient handling of this class of nets. We have first established a polynomial-time algorithm to check if a given SPN is a PF-SPN. This is an interesting result, in contrast with the general computation of T-semiflows which may produce an exponential number of T-semiflows (with respect to the size of the net). Then, we have proven a rate independent structural characterization of PF-SPN, which can also be checked in polynomial time. We call $\overline{\Pi}$-nets the subclass of Π-nets satisfying this criterion. Moreover, for $\overline{\Pi}$-nets, we are able to define transition rates globally dependent of components of the net "not related with" the considered transition, so that we can model complex dependency of activities on some other ones. Third, we have investigated untimed properties for the class of PF-SPN. We have shown that Π-nets, and even $\overline{\Pi}$-nets do not fit in any standard class of PN. Nevertheless, we have proved specific properties for deadlock-freeness, liveness and reverse nets for Π-nets. For what concerns liveness/reachability in Π-nets and $\overline{\Pi}$-nets, we were able to somewhat refine complexity bounds known for general PN. Finally, with examples and one proposition, we have given some answers, both positive and negative, to the problem of potential reachability, i.e. reachability based upon structural properties of the net. The interested reader will find detailed proofs and full versions of results in [10].

References

1. G. Balbo, S. C. Bruell, and M. Sereno. Arrival theorems for product-form stochastic Petri nets. In *Proc. 1994 ACM SIGMETRICS Conference*, pages 87–97, Nashville, Tennessee, USA, May 1994. ACM.
2. R. J. Boucherie and M. Sereno. On closed support t-invariants and traffic equations. *Journal of Applied Probability*, (35):473–481, 1998.
3. S. C. Bruell and G. Balbo. *Computational Algorithms for Closed Queueing Networks*. Elsevier North-Holland, New York, 1980.
4. J. P. Buzen. Computational algorithms for closed queueing networks with exponential servers. *Communications of the ACM*, 16(9):527–531, September 1973.
5. E. Cardoza, R.J. Lipton, and A.R. Meyer. Exponential space complete problems for Petri nets and commutative semigroups. In *Proc. of the 8th Annual Symposium on Theory of Computing*, pages 50–54, 1976.
6. A. Cheng, J. Esparza, and J. Palberg. Complexity results for 1-safe nets. In *Proc. of the 13th Conference on Foundations of Software Technology and Theoretical Computer Science*, Bombay, India, 1993.
7. J. L. Coleman, W. Henderson, and P. G. Taylor. Product form equilibrium distributions and an algorithm for classes of batch movement queueing networks and stochastic Petri nets. *Performance Evaluation*, 26(3):159–180, September 1996.
8. J. Esparza. *Decidability and Complexity of Petri nets problems – an introduction*, pages 374–428. 1998. in [22].

9. J. Esparza and M. Nielsen. Decidability issues for Petri nets – a survey. *Journal of Information Processing and Cybernetics*, 30(3):143–160, 1994. former version in Bulletin of the EATCS, volume 52, pages 245–262, 1994.
10. S. Haddad, P. Moreaux, M. Sereno, and M. Silva. Revisiting Product Form Stochastic Petri Nets. Technical report, LERI-*RESYCOM*, Université de Reims Champagne-Ardenne, Reims, France, June 2001.
11. W. Henderson, D. Lucic, and P.G. Taylor. A net level performance analysis of stochastic Petri nets. *Journal of Australian Mathematical Soc. Ser. B*, 31:176–187, 1989.
12. W. Henderson and P.G. Taylor. Aggregation methods in exact performance analysis of stochastic Petri nets. In *Proc. 3^{rd} Intern. Workshop on Petri Nets and Performance Models*, pages 12–18, Kyoto, Japan, December 1989. IEEE-CS Press.
13. W. Henderson and P.G. Taylor. Embedded processes in stochastic Petri nets. *IEEE Transactions on Software Engineering*, 17:108–116, February 1991.
14. N.D. Jones, L.H. Landweber, and Y.E. Lien. Complexity of some problems in Petri nets. *Theoretical Computer Science*, 4:277–299, 1977.
15. A. A. Lazar and T. G. Robertazzi. Markovian Petri net protocols with product form solution. In *Proc. of INFOCOM '87*, pages 1054–1062, San Francisco, CA, USA, 1987.
16. R.J. Lipton. The reachability problem requires exponential space. Technical report, Dpt. of Computer Science, Yale University, 1976.
17. J. Martinez and M. Silva. A simple and fast algorithm to obtain all invariants of a generalized Petri net. In *Proc. 2^{nd} European Workshop on Application and Theory of Petri Nets*, Bad Honnef, West Germany, September 1981. Springer Verlag.
18. E.W. Mayr and A.R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, (46):305–329, 1982.
19. T. Murata. Petri nets: properties, analysis, and applications. *Proceedings of the IEEE*, 77(4):541–580, April 1989.
20. J. L. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1981.
21. M. Reiser and S. S. Lavenberg. Mean value analysis of closed multichain queueing networks. *Journal of the ACM*, 27(2):313–322, April 1980.
22. W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets I: Basic models*. Number 1491 in LNCS. Springer–Verlag, June 1998. Advances in Petri nets.
23. M. Sereno and G. Balbo. Computational algorithms for product form solution stochastic Petri nets. In *Proc. 5^{th} Intern. Workshop on Petri Nets and Performance Models*, pages 98–107, Toulouse, France, October 1993. IEEE-CS Press.
24. M. Sereno and G. Balbo. Mean value analysis of stochastic Petri nets. *Performance Evaluation*, 29(1):35–62, 1997.
25. K. C. Sevcik and I. Mitrani. The distribution of queueing network states at input and output instants. *Journal of the ACM*, 28(2):358–371, April 1981.
26. M. Silva. *Las Redes de Petri en la Automatica y la Informatica*. Ed. AC, Madrid, Spain, 1985. In Spanish.
27. M. Silva, E. Teruel, and J.M Colom. *Linear algebraic and linear programming techniques for the analysis of Place/Transition net systems*, pages 309–372. 1998. in [22].