ARTICLE IN PRESS

Product-form and stochastic Petri nets: a structural approach

S. Haddad[a], P. Moreaux[a,b,*], M. Sereno[c], M. Silva[d]

[a] *LAMSADE, Université Paris Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16, France*
[b] *CReSTIC, Université de Reims Champagne-Ardenne, Dpt. Mathématiques et Informatique, BP 1039, 51687 Reims Cedex 02, France*
[c] *Dipartimento di Informatica, Università di Torino, Corso Svizzera 185, 10149 Torino, Italy*
[d] *Universidad de Zaragoza, Dep. de Ingenieria Elec. e Informatica, Maria de Luna 3, E50015 Zaragoza, Spain*

Abstract

Stochastic Petri nets (SPNs) with product-form solution are nets for which there is an analytic expression of the steady-state probabilities with respect to place markings, as it is the case for product-form queueing networks with respect to queue lengths. The most general kind of SPNs with product-form solution introduced by Coleman et al. (and denoted here by $S\Pi$-nets) suffers a serious drawback: the existence of such a solution depends on the values of the transition rates. Thus since their introduction, it is an open question to characterize $S\Pi$-nets with product-form solution for any values of the rates. A partial characterization has been obtained by Henderson et al. However, this characterization does not hold for every initial marking and it is expressed in terms of the reachability graph. In this paper, we obtain a purely structural characterization of $S\Pi$-nets for which a product-form solution exists for any value of probabilistic parameters of the SPN and for any initial marking. This structural characterization leads to the definition of $S\Pi^2$-nets (Stochastic Parametric Product-form Petri nets). We also design a polynomial time (with respect to the size of the net structure) algorithm to check whether a SPN is a $S\Pi^2$-net. Then, we study qualitative properties of Π-nets and Π^2-nets, the non-stochastic versions of $S\Pi$-nets and $S\Pi^2$-nets: we establish two results on the complexity bounds for the liveness and the reachability problems, which are central problems in Petri nets theory. This set of results complements previous studies on these classes of nets and improves the applicability of product-form solutions for SPNs.
© 2004 Elsevier B.V. All rights reserved.

Keywords: Performance evaluation; Stochastic Petri net; Product-form; Subclasses of Petri nets

* Corresponding author. Tel.: +33 326 913 397; fax: +33 326 913 344.
E-mail addresses: haddad@lamsade.dauphine.fr (S. Haddad), moreaux@lamsade.dauphine.fr (P. Moreaux), matteo@di.unito.it (M. Sereno), silva@posta.unizar.es (M. Silva).

ARTICLE IN PRESS

1. Introduction

Stochastic models of discrete events systems have been proven for many years to be powerful tools for modelling and evaluating the performances of systems like parallel and distributed systems, database systems, communication networks, etc. For the steady-state performance analysis, it is usually necessary to compute the steady-state distribution of a Continuous Time Markov Chain (CTMC) derived from the model. This is the case for Queueing Networks (QN) and Stochastic Petri Nets (SPNs) models. Because of the huge state space describing these complex systems, it is often extremely difficult to compute the exact numerical solution of this CTMC. This situation is obviously even worse with infinite state space models which prevent such a computation. A first attempt to overcome this so-called state space explosion problem is to leave the domain of exact solutions. Three main approaches have been and are still developed in this area: discrete-event simulation, approximate methods and computation of bounds. If we wish to stay in the framework of exact methods, then we need to 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 the mathematical problem. The latter approach, which is the framework of this paper, aims at describing the steady-state probabilities and other performance measures as *functions* of a fixed set of parameters of the states, derived from the model structure. Models for which such solutions may be developed are said product-form models, since the structure of the functions are usually a product of elementary terms corresponding to the parameters. The analysis of Queueing Networks (QNs) with product-form (PF-QN) [25,24] provided the first important results in this direction. The general approach of product-form analysis for QNs is made up of four main parts. First it is necessary to recognize, among all possible QNs, the properties that ensure a PF-solution. This has led to well-known classifications of QNs (closed, open or mixed, single or multi-classes, disciplines of the queues, etc. [24,18,2]) together with the introduction of several notions such as visit-ratios, routing, etc. and required relations among them. Then, the visit-ratios, related to the load of customers in each station are computed using the parameters of these stations and the routes followed by the customers (the routing process). From the structure of the QN and the visit-ratios, it has been shown, for several classes of QNs, that the steady-state probabilities of the model may be expressed as a product of terms, depending on the state of each station, up to a normalization constant. Finally, the normalization constant must be determined and performance indices (throughput, length and service times, utilization, etc.) can be both derived by adapted algorithms such as, for instance, the Mean Value Analysis (MVA) method [7] and the convolution algorithm [40].

Due to the explicit modelling of competition and concurrency, the stochastic Petri nets model [1], is an attractive (and complementary with respect to QNs) modelling paradigm when studying performance of systems for which these complex phenomena have an important impact on their behaviour. In contrast with PF-QN, it is almost always necessary to turn towards approximate methods (see [4] for a presentation of various methods, and [3] for synchronized QN) to deal with such systems, although analogous trends can also be observed with SPNs (see for example [8]). From the late 1980's, Product-Form Stochastic Petri Net (PF-SPN) were introduced as an attempt to cope with the state explosion problem for SPN models. Although the global steps for PF-SPNs and PF-QNs are analogous, specific problems arise with PF-SPNs. The first one, and in some sense, the most important one, is to find what properties of SPN are relevant with regard to the existence of a product-form solution (PF-solution). Since SPN models take into account complex synchronization schemas, it is not surprising that the search for required properties to ensure a PF-solution gave rise to many proposals. Historically, we observe that works started from purely

behavioural properties (i.e. by a analysis of the reachability graph) as in [30], and then progressively introduced more and more structural parameters to ensure a PF-solution [31,17,19,5]. In [5], Boucherie introduces the product of Competing Markov Chains (PPCMC) and shows that several SPN generate such chains. In his paper, the author identifies some structural properties of the SPN, but they are not studied for their own sake and the PF-solution are established on a case by case basis. The work [19] introduces several important structural properties: the identification of cyclic activities in the net (T-flows, i.e. sequences of transitions which may be repetitively fired while leaving the markings unchanged) and the existence of "balanced" input and output bags of transitions, providing structural conditions for shared resources usage. The importance of T-flows (more precisely, closed support minimal T-semiflows) was emphasized with the identification of the class of so-called Π-nets [6] which is now the starting point for obtaining SPNs with a PF-solution. The crucial breakthrough was done by Henderson et al. when they established a numerical condition on the parameters of the Π-net which ensures the existence of a PF-solution [19]. Unfortunately, this technical numerical condition has no intuitive interpretation in relation with the modelling, and above all, it relies the existence of a PF-solution on the *numerical values* of the rates of the transitions of the net, in contrast with other models such as QNs. Thus in a later work [13], these authors have characterized Π-nets for which the existence of the product-form solution does not depend on the particular values of the nets. However, their characterization suffers two drawbacks: on the one hand, it holds only when the initial marking is sufficiently large and, on the other hand, this condition is expressed on the reachability graph of the net. Since the aim of the product-form methods is to avoid the construction of such a graph, checking this condition does not make sense.

Thus the existence of a structural characterization of Π-nets with PF-solution whatever are the rates (i.e. expressed with respect to the structure of the net) was still an open problem. The main contribution of the present work is *such a structural characterization*. Furthermore, we show that this characterization can be checked in polynomial time with respect to the size of the net. In the sequel, we will denote such nets as rate-insensitive PF-Π-nets.

In order to informally explain our characterization, we first recall what is a Π-net while simultaneously giving an intuitive interpretation. In a Π-net, the transitions are partitioned in components. The transitions of each component model the activities of a group of (virtual) clients *not explicitly represented in the net*. Each place represents a kind of resource and the input (resp. the output) arcs of a transition represent the resources consumed (resp. produced) by the activity associated to the transition. In Π-nets, each multi-set of resources consumed (resp. produced) by a transition of a component must be exactly the multi-set of resources produced (resp. consumed) by another transition of the same component. In other words, the partition of transitions may be deduced from these multi-sets (called in the sequel *input/output bags*) by requiring that an input/output bag belongs to a single component. The concurrent accesses to resources between two transitions belonging to different components are not restricted in any way. The product-form solution we look for should be a product of factors over the components multiplied by a normalization constant. The factor associated to a component should be again a product of sub-factors over the "states" of the clients associated to a component. Here is the crucial point for the existence of the product-form solution. How can we characterize the number of clients in a given state or at least the variation of this number? Starting from the above interpretation, the entry into a state is witnessed by the production of the associated input/output bag due to some transition firing in the component and the state exit is witnessed by the production of the input/output bag due to some other transition firing. In other words given a marking **m** reached by a sequence σ, the variation of the number of clients in a "state" should be given by a weighted sum of transition occurrences in σ (i.e. an item of *the count vector*

defined by [13]). Consequently in order to obtain a product-form solution, the count vector should be a function of the reachable marking. This observation is the core of the partial characterization given in [13]. Actually, what we will prove here is that a Π-net is a rate-insensitive PF-Π-net iff *every item of the count vector is obtained by a linear mapping applied on the difference between the current marking and the initial marking*. Then in such nets, a (virtual) client state is characterized by a linear combination of places whose current marking gives the variation of the number of virtual clients in this state. From an algebraic point of view, these linear combinations are partial flows and can be seen as an extension of the synchronic distance relation [46] which quantifies the firing dependencies between transitions. We call Π^2-nets, the Π-nets for which such algebraic relations hold, and $S\Pi^2$-nets (Stochastic Parametric Product-form PNs) their stochastic version.

Furthermore, since in $S\Pi^2$-nets we have precisely identified virtual clients states, we also add to the functional dependencies of the rates of the transitions in a component, a uniform dependency based on the states of the clients in the other components. We illustrate the interest of this extension on a short example. We also design an algorithm for the verification of our characterization whose time complexity is polynomial with respect to size of the net.

The second group of results presented here relates to the complexity of the reachability and the liveness problems for Π-nets and Π^2-nets. Several standard subclasses of nets have been identified since the introduction of PNs and many important complexity results have been obtained for these classes. It was shown [43] that Π-nets cannot be directly classified with respect to the standard subclasses of PNs. This motivated the study of the complexity of these central problems. In this respect, we provide different complexity bounds with tight lower and upper bounds in two cases.

The organization of the paper is as follows. In Section 2, we remind Π-nets for which we also provide a new and efficient membership algorithm. Then we present in Section 3 our main contribution: the structural characterization of rate-insensitive PF-Π-nets, i.e. the definition of $S\Pi^2$-nets and the proof that a Π-net is a rate-insensitive PF-Π-net iff it is a $S\Pi^2$-net. In Section 4, we study the complexity of the liveness and the reachability problems in Π-nets and Π^2-nets. We conclude with a review of open problems and future works in Section 5.

2. Π-nets

In this section, we summarize previous results about SPN with PF-solution. One may find introductory presentations of Petri net concepts for instance in [36,38,45] and we remind the reader only with definitions necessary to understand product-form results for stochastic Petri nets.

A *Petri Net system* (or PN for short) is a tuple $\mathcal{S} = (\mathcal{P}, \mathcal{T}, W, \mathbf{m}_0)$, where $\mathcal{P}$ and $\mathcal{T}$ are disjoint sets of *places* and *transitions* (with $|\mathcal{P}| = n_p$ and $|\mathcal{T}| = n_t$); $W : \mathcal{P} \times \mathcal{T} \cup \mathcal{T} \times \mathcal{P} \to \mathbb{N}$ 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 $W(i, j)$); $\mathbf{m}_0$ is the initial marking. We denote by $\mathcal{N} = (\mathcal{P}, \mathcal{T}, W)$ the Petri net, derived from $\mathcal{S}$ without considering the initial marking.

For a given transition $t_j \in \mathcal{T}$, its *preset* and *postset* are given by $^\bullet t_j = \{p_i | W(i, j) > 0\}$ and $t_j^\bullet = \{p_i | W(j, i) > 0\}$, respectively. In the same manner we can define the *preset* and *postset* of a given place. We also define the *input vector* $\mathbf{i}(t_j) = [W(l, j), W(2, j), \ldots, W(n_p, j)]$ and the *output vector* $\mathbf{o}(t_j) = [W(j, 1), W(j, 2), \ldots, W(j, n_p)]$ of a transition t_j. The matrix $\mathbf{C}$ with entries $\mathbf{C}[i, j] = W(j, i) - W(i, j)$ is called the *incidence matrix* of $\mathcal{N}$.

ARTICLE IN PRESS

Fig. 1. An introductory example of Π-net.

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 $\mathbf{m}' = \mathbf{m} - \mathbf{i}(t_j) + \mathbf{o}(t_j) = \mathbf{m} + \mathbf{C}[\cdot, j]$ ($\mathbf{C}[\cdot, j]$ is the jth column of $\mathbf{C}$), which is denoted by $\mathbf{m} \rightarrow^{t_j} \mathbf{m}'$. The set of all the markings reachable from $\mathbf{m}_0$ is called the *reachability set* of $\mathcal{S}$, and is denoted by RS($\mathbf{m_0}$).

Semiflows are non-null natural annullers of $\mathbf{C}$. Right and left annullers are called T- and P-semiflows, respectively. A semiflow $\mathbf{s}$ is called *minimal* when its support (i.e., the set $\|\mathbf{s}\|$ of places or transitions corresponding to the non-zero components of the vector $\mathbf{s}$) is not a proper superset of the support of any other semiflow, and the g.c.d. of its elements is 1.

2.1. Introductory example and definitions

Obviously, general SPNs may not allow a PF-solution. Even before looking at the stochastic problems, we must enforce some structure in their behavior, as for QN. Moreover, as usual with PNs, we want to deduce these structures from the syntax of the net and not by examining its reachability graph. Let us give an example of such a net that we will use throughout the paper. The net of Fig. 1 models a system with two groups of concurrent activities, for instance computation tasks and interactive tasks performed by two categories of "clients" (batch jobs and human beings). In PN, basic activities are modelled by *transitions*; we have two sets of transitions: t_1, t_2, t_3 (batch jobs) and t_4, t_5, t_6, t_7 (interactive tasks). The cyclic behaviour of activities is allowed by the balance between the input/output bags of transitions belonging to the same component. Inside a component, a client may have several behaviours like (t_4, t_5) or (t_4, t_6, t_7) for interactive clients. Activities use or produce resources, may be in a competitive way. In PN models, (passive) resources are modelled by tokens in places and input/output bags of transitions indicate how each activity manages these resources. For instance, p_1 could represent free CPUs, required by clients of both components, while p_2, p_3 model specific resources for batch jobs only like disks.

Such a structure in PN models is captured by the following definitions.

Definition 1. A subset $\mathcal{T}'$ of transitions is said to be closed if $\bigcup_{t \in \mathcal{T}'} \mathbf{i}(t) = \bigcup_{t \in \mathcal{T}'} \mathbf{o}(t)$.

ARTICLE IN PRESS

We will denote by $\mathcal{R}(\mathcal{T}') = \bigcup_{t \in \mathcal{T}'}\{\mathbf{i}(t), \mathbf{o}(t)\}$ the set of input and output bags for transitions in $\mathcal{T}'$.

We can now recall the definition of the Π-net class of PN. All Stochastic Petri Nets with PF-solution studied in this paper, will be Π-nets.

Definition 2. $\mathcal{N}$ is a Π-net if $\forall t \in \mathcal{T}$ there exists a minimal T-semiflow $\mathbf{x}$ such that $t \in \|\mathbf{x}\|$, 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. Finally, the structural category corresponding to a group of activities is defined as follows.

Definition 3. Two different minimal closed support T-semiflows $\mathbf{x}$ and $\mathbf{x}'$ are said to be freely related [9] (denoted as $(\mathbf{x}, \mathbf{x}') \in \mathrm{FR}$), if there exists $t \in \|\mathbf{x}\|$ and $t' \in \|\mathbf{x}'\|$ such that $\mathbf{i}(t) = \mathbf{i}(t')$.

FR^* is the transitive closure of FR. It induces an equivalence relation among transitions (and among their input/output bags): $(t, t') \in \mathrm{FR}^*$ iff $t \in \|\mathbf{x}\|$, $t' \in \|\mathbf{x}'\|$ and $(\mathbf{x}, \mathbf{x}') \in \mathrm{FR}^*$. We denote by $\mathcal{C}(t)$ the FR^* class of t or the FR^* class of its input/output bags (since no confusion can arise).

In the example of Fig. 1, we can see that there are three minimal T-semiflows $\mathbf{x_1} = [1, 1, 1, 0, 0, 0, 0]$, $\mathbf{x_2} = [0, 0, 0, 1, 1, 0, 0]$ and $\mathbf{x_3} = [0, 0, 0, 1, 0, 1, 1]$, with $\|\mathbf{x_1}\| = \{t_1, t_2, t_3\}$, $\|\mathbf{x_2}\| = \{t_4, t_5\}$ and $\|\mathbf{x_3}\| = \{t_4, t_6, t_7\}$. We observe that $\bigcup_{t_j \in \|\mathbf{x_1}\|} \mathbf{i}(t_j) = \{[1, 0, 0, 0, 0, 0], [0, 1, 0, 0, 0, 0], [0, 0, 1, 0, 0, 0]\} = \bigcup_{t_j \in \|\mathbf{x_1}\|} \mathbf{o}(t_j)$, $\bigcup_{t_j \in \|\mathbf{x_2}\|} \mathbf{i}(t_j) = \{[1, 0, 0, 1, 0, 0], [0, 0, 0, 0, 1, 0]\} = \bigcup_{t_j \in \|\mathbf{x_2}\|} \mathbf{o}(t_j)$ and $\bigcup_{t_j \in \|\mathbf{x_3}\|} \mathbf{i}(t_j) = \{[1, 0, 0, 1, 0, 0], [0, 0, 0, 0, 1, 0], [0, 0, 0, 0, 0, 1]\} = \bigcup_{t_j \in \|\mathbf{x_3}\|} \mathbf{o}(t_j)$. The three T-semiflows have closed support set. Since any transition belongs to a closed support minimal T-semiflow, this net is a Π-net. We can also note that we have six input/output bags $\mathbf{i}(t)$ and two FR^* classes, $\mathcal{C}_1 = \{t_1, t_2, t_3\}$ and $\mathcal{C}_2 = \{t_4, t_5, t_6, t_7\}$ since $\mathbf{i}(t_5) = \mathbf{i}(t_6)$.

2.2. Membership problem

From the definition of Π-nets, we can decide whether 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 [33]. We present now an algorithm that allows to recognize whether a net is a Π-net *in polynomial time* (unless explicitly mentioned, all complexity results in the paper are with respect to the size of the net, i.e. the number of places, transitions, arcs and the binary representation of valuations). The soundness of the algorithm is based on the following lemma.

Lemma 4. *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 ith component of $\mathbf{x}$);*
(ii) *$\|\mathbf{x}\|$ may be ordered as $\{t_{j_0}, t_{j_1}, \ldots, t_{j_{h-1}}\}$ such that $\mathbf{o}(t_{j_i}) = \mathbf{i}(t_{j_{i+1 \bmod h}})$ (for $i = 0, 1, \ldots, h - 1$), and $l \neq l' \Rightarrow \mathbf{i}(t_{j_l}) \neq \mathbf{i}(t_{j_{l'}})$.*

Proof. Let $\mathbf{x}$ a closed support minimal T-semiflow and $t_{j_0} \in \|\mathbf{x}\|$, then $\exists t_{j_1} \in \|\mathbf{x}\|$ such that $\mathbf{i}(t_{j_1}) = \mathbf{o}(t_{j_0})$. We iterate the procedure until we find a transition $t_{j_k} \in \|\mathbf{x}\|$ such that $\exists k' < k$ with $\mathbf{i}(t_{j_k}) = \mathbf{i}(t_{j_{k'}})$. We get a new closed T-semiflow $\mathbf{x}'$ with $\|\mathbf{x}'\| = \{t_{j_{k'}}, \ldots, t_{j_{k-1}}\}$. By construction $\mathbf{x}'$ has all the required properties, i.e., $\mathbf{o}(t_{j_i}) = \mathbf{i}(t_{j_{i+1 \bmod (k-k')}})$ (for $i = k', \ldots, k - 1$), $l \neq l' \Rightarrow \mathbf{i}(t_{j_l}) \neq \mathbf{i}(t_{j_{l'}})$, and $\|\mathbf{x}'\| \subseteq \|\mathbf{x}\|$. The minimality of $\mathbf{x}$ implies that $\|\mathbf{x}'\| = \|\mathbf{x}\|$ (another proof is given in [6]). $\qquad\square$

2.2.1. 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 $\mathbf{o}(t_{j_i}) = \mathbf{i}(t_{j_{i+1 \bmod h}})$ (for $i = 0, 1, \ldots, h-1$). The algorithm *Verify Π-net* exploits this feature for checking if a net is a Π-net.

Algorithm (Verify Π-net).

begin
$\mathcal{L} \leftarrow \mathcal{T}$
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 */
end

We point out that the algorithm yields a covering set of closed support minimal T-semiflows (if the PN is a Π-net).

We can easily see that any transition is analyzed exactly once during the execution of the algorithm. Moreover, other tests may require at most $O(n_p \times n_t)$ elementary computations, so that the complexity of the algorithm that allows to recognize if a given net satisfy Definition 2 is at most $O(n_t^2 \times n_p)$ (i.e. $O(|\mathcal{T}|^2 \times |\mathcal{P}|)$).

2.3. From Π-nets to PF-SΠ-nets

Stochastic Petri Net systems (or SPN for short) are PN where the transitions have exponentially distributed firing delays with rate μ_t. The sojourn time in the marking $\mathbf{m}$ before the firing of t is exponentially distributed with rate $\mu(t, \mathbf{m})$ (if $\mu(t, \mathbf{m}) = \mu_t$ we say that the firing rates are marking independent). In the rest of the paper, we assume that the CTMC underlying the SPN (with state space RS($\mathbf{m}_0$)) is ergodic. Since the initial marking of a Π-net is a home state (due to the existence of the closed T-semiflows) then

the ergodicity is ensured if and only if the invariant measure associated to the product-form solution is finite. Several authors added stochastic characteristics to Π-nets (Stochastic Π-nets, $S\Pi$-nets), leading to PF-SPN under specific conditions.

Following [19,22], we consider $S\Pi$-nets, with rates $\mu(t, \mathbf{m})$ satisfying, for t enabled in $\mathbf{m}$, the relation

$$\mu(t, \mathbf{m}) = \mu_t \frac{\psi(\mathbf{m} - \mathbf{i}(t))}{\phi(\mathbf{m})}, \quad \text{where } \psi \geq 0, \ \phi > 0. \tag{1}$$

The functions ψ and ϕ can be thought of as "potential functions", the state dependent firing rate of transition t in $\mathbf{m}$ being the product of its intrinsic firing rate μ_t and the ratio of the functions ψ and ϕ evaluated at the states that exist after and before consuming tokens, respectively. Marking independent firing rates can be modelled by choosing $\psi = \phi = 1$.

In order to adopt a virtual client perspective, we define $\forall \mathbf{i} \in \mathcal{R}(\mathcal{T})$, $\mu(\mathbf{i}) = \sum_{t \in \mathcal{T}, \mathbf{i}(t) = \mathbf{i}} \mu_t$ the intrinsic exit rate from the state associated to $\mathbf{i}$ and $\forall t$ s.t. $\mathbf{i}(t) = \mathbf{i}$, $\mathbf{P}[\mathbf{i}, t] = \mu_t / \mu(\mathbf{i})$ the probability to fire t when exiting this state. Starting from the client point of view, we can also reformulate the firing rate of a transition:

$$\mu(t, \mathbf{m}) = \mu(\mathbf{i}(t)) \frac{\psi(\mathbf{m} - \mathbf{i}(t))}{\phi(\mathbf{m})} \mathbf{P}[\mathbf{i}(t), t]. \tag{2}$$

A service function of the form (2) is common in the literature on PF-QNs (see, for example [20] or [44]). In the SPN context, it was first used in [19] where several examples are given which illustrate its range of application. For single movement queueing networks, a form with $\psi = \phi$ first appeared in [28], as a generalization of the state dependent service rates introduced in [24].

As for QNs, PF-solutions for SPNs are based on the analysis of underlying Markov chains (MCs). It is then convenient to study an auxiliary Discrete Time Markov Chain (DTMC) y, called the *routing process* [19] of the $S\Pi$-net, with states being the input/output bags. Let us define its transition matrix:

$$\mathbf{P}[\mathbf{i}, \mathbf{i}'] \stackrel{\text{def}}{=} \sum_{\substack{t \in \mathcal{T} \\ \mathbf{i}(t) = \mathbf{i}, \mathbf{o}(t) = \mathbf{i}'}} \mathbf{P}[\mathbf{i}(t), t].$$

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

$$\forall t \in \mathcal{T}, \quad v(\mathbf{i}(t)) = \sum_{t' \in \mathcal{T}} v(\mathbf{i}(t')) \mathbf{P}[\mathbf{i}(t'), \mathbf{i}(t)], \tag{3}$$

or equivalently,

$$\forall \mathbf{i} \in \mathcal{R}(\mathcal{T}), \quad v(\mathbf{i}) = \sum_{\mathbf{i}' \in \mathcal{R}(\mathcal{T})} v(\mathbf{i}') \mathbf{P}[\mathbf{i}', \mathbf{i}]. \tag{4}$$

Boucherie and Sereno [6] showed that traffic equations and structural properties of a net are closely related.

Theorem 5 (From [6]). *Let $\mathcal{C}$ be the set of input/output bags classes of a $S\Pi$-net with respect to the relation* FR*. *Then*

(i) $\mathcal{C}$ *is a partition of the routing chain y into* $|\mathcal{C}|$ *irreducible absorbing sub-chains on each* $C \in \mathcal{C}$.
(ii) *the traffic* equation (4) *are equivalent to the* $|\mathcal{C}|$ *systems of equations*

$$\forall \mathbf{i} \in C, \quad v(\mathbf{i}) = \sum_{\mathbf{i}' \in C} v(\mathbf{i}')\mathbf{P}[\mathbf{i}', \mathbf{i}], \tag{5}$$

which are independent.
(iii) *each system* (5) *admits a unique positive solution up to a multiplicative constant.*

Unfortunately, the existence of a positive solution for the Traffic equation (3) is not a sufficient condition to assert a PF-solution for $S\Pi$-nets. The following result from Coleman et al. [14], 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 $\mathbf{w}_f = [w_1, \ldots, w_{n_t}]$ as

$$\mathbf{w}_f = \left[\log\left(\frac{f(\mathbf{i}(t_1))}{f(\mathbf{o}(t_1))}\right), \log\left(\frac{f(\mathbf{i}(t_2))}{f(\mathbf{o}(t_2))}\right), \ldots, \log\left(\frac{f(\mathbf{i}(t_{n_t}))}{f(\mathbf{o}(t_{n_t}))}\right) \right]. \tag{6}$$

There are many such functions *f* corresponding to different solutions of 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 6 (Product-Form for SPN (from [14])). *Let* $f = v/\mu$ *with* v *a solution for the traffic equations. The equilibrium distribution for the SPN has the form*

$$\pi(\mathbf{m}) = \frac{1}{G}\phi(\mathbf{m}) \prod_{i=1}^{n_p} y_i^{m_i}, \quad \forall \mathbf{m} \in RS(\mathbf{m_0}) \tag{7}$$

if and only if $\mathrm{Rank}(C) = \mathrm{Rank}([C|\mathbf{w}_f])$, *where* $[C|\mathbf{w}_f]$ *is the matrix* $\mathbf{C}$ *augmented with the row* $\mathbf{w}_f$ *and* G *is a normalization constant.*

In this case, the n_p-*component vector* $\mathbf{l} = [\log(y_1), \ldots, \log(y_{n_p})]$, *satisfies the matrix equation* $-\mathbf{l} \cdot C = \mathbf{w}_f$.

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

2.4. Examples of SΠ-nets

Let us work out two detailed examples of PF-$S\Pi$-nets. The first one complements the study of the introductory example, and the second one shows a more complex situation related to the rank condition of Theorem 6.

2.4.1. Example 1

In this example we briefly review the procedure used to obtain the equilibrium distribution for the $S\Pi$-net of Fig. 1. For additional details the reader is referred to [6,14,19,21,22]. Since we know that the net is a $S\Pi$-net, and $\mathcal{C} = \{C_1, C_2\}$ with $C_1 = \{t_1, t_2, t_3\}$ and $C_2 = \{t_4, t_5, t_6, t_7\}$, there is a solution for each of the systems (5):

- class C_1:

$$v(\mathbf{i}(t_1)) = v(\mathbf{i}(t_2)),$$

$$v(\mathbf{i}(t_2)) = v(\mathbf{i}(t_3)),$$

$$v(\mathbf{i}(t_3)) = v(\mathbf{i}(t_1));$$

- class C_2:

$$v(\mathbf{i}(t_4)) = \frac{\mu_5}{\mu_5 + \mu_6} v(\mathbf{i}(t_5)) + v(\mathbf{i}(t_7)),$$

$$v(\mathbf{i}(t_5)) = v(\mathbf{i}(t_4)) \, (\text{and } \mathbf{i}(t_6) = \mathbf{i}(t_5)),$$

$$v(\mathbf{i}(t_7)) = \frac{\mu_6}{\mu_5 + \mu_6} \times v(\mathbf{i}(t_5)).$$

Obviously, each system has a redundant equation. Setting for instance $v(\mathbf{i}(t_1)) = 1$ and $v(\mathbf{i}(t_4)) = 1$, we get $v(\mathbf{i}(t_1)) = v(\mathbf{i}(t_2)) = v(\mathbf{i}(t_3)) = v(\mathbf{i}(t_4)) = v(\mathbf{i}(t_5)) = 1$, and $v(\mathbf{i}(t_7)) = \mu_6/(\mu_5 + \mu_6)$, from which we obtain $f(\mathbf{i}(t_1)) = 1/\mu_1$, $f(\mathbf{i}(t_2)) = 1/\mu_2$, $f(\mathbf{i}(t_3)) = 1/\mu_3$, $f(\mathbf{i}(t_4)) = 1/\mu_4$, $f(\mathbf{i}(t_5)) = 1/\mu_5$ and $f(\mathbf{i}(t_7)) = 1/(\mu_5 + \mu_6)$. The row vector $\mathbf{w}_f$ is

$$\mathbf{w}_f = \left[\log\left(\frac{\mu_2}{\mu_1}\right), \log\left(\frac{\mu_3}{\mu_2}\right), \log\left(\frac{\mu_1}{\mu_3}\right), \log\left(\frac{\mu_5}{\mu_4}\right), \log\left(\frac{\mu_4}{\mu_5}\right), \log\left(\frac{(\mu_5 + \mu_6)\mu_7}{\mu_5\mu_6}\right), \right.$$
$$\left. \log\left(\frac{\mu_4\mu_6}{(\mu_5 + \mu_6)\mu_7}\right) \right].$$

It can be verified that the rank condition $\text{rank}(\mathbf{C}) = \text{rank}([\mathbf{C}|\mathbf{w}_f])$ is satisfied independently of the rate values and a simple derivation gives the equilibrium distribution of the introductory example:

$$\pi(\mathbf{m}) = \frac{1}{G} \left(\frac{\mu_1}{\mu_2}\right)^{\mathbf{m}[p_2]} \left(\frac{\mu_2}{\mu_3}\right)^{\mathbf{m}[p_3]} \left(\frac{\mu_4}{\mu_5}\right)^{\mathbf{m}[p_5]} \left(\frac{\mu_6}{\mu_5 + \mu_6}\frac{\mu_4}{\mu_7}\right)^{\mathbf{m}[p_6]}.$$

2.4.2. Example 2

The SPN shown in Fig. 2, taken form [14], represents an SPN in which the rank condition is not satisfied independently of the rate values. 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 $f(\mathbf{i}(t_i)) = 1/\mu_i$ for $i = 1, \ldots, 4$. The vector $\mathbf{w}_f$ is given by

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

Fig. 2. A rate-sensitive PF-SPN (example 2).

The rank conditions are $w_2 + 2w_1 = 0$, $w_3 - 2w_1 = 0$, and $w_1 + w_4 = 0$, which implies, $f(\mathbf{i}(t_2))/f(\mathbf{i}(t_3))(f(\mathbf{i}(t_1))/f(\mathbf{i}(t_4)))^2 = 1$, $f(\mathbf{i}(t_3))/f(\mathbf{i}(t_2))(f(\mathbf{i}(t_4))/f(\mathbf{i}(t_1)))^2 = 1$, and $\mathbf{l} = 1$ 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 $\mu_2/\mu_3 = (\mu_4/\mu_1)^2$. If this condition is met, Theorem 6 applies, and, letting $y_2 = 1$ gives $y_1 = f(\mathbf{i}(t_1))/f(\mathbf{i}(t_4))$. Finally, $\pi_f(\mathbf{m}) = [\mu_4/\mu_1]^{\mathbf{m}[p_1]}$.

The major drawbacks of the above results about a PF-solution for $S\Pi$-nets are threefold. First, the technical rank condition (Theorem 6) has no intuitive interpretation in relation with the modelling. Second, and this is, for sure the most important point, the existence of a PF-solution depends on the *numerical values* of the rates of the transitions of the net, in contrast with other models. Note also that the log functions involved in the formulae lead to numerical sensitivity of the verification of the condition. Finally, there is an apparent contradiction with the negligible theoretical occurrence of the condition satisfaction and the fact that it occurs quite often in practice. These remarks motivated the search for a structural characterization of $S\Pi$-nets with PF-solution, which is exposed in the next section.

3. Π^2-nets: definition and performance analysis

In this section we define the class of $S\Pi^2$-nets (Stochastic Parametric Product-form Nets) (and Π^2-nets, their non-stochastic version) and we will show that it exactly corresponds to the class of rate-insensitive PF-Π-nets. Moreover, we introduce a more general dependency of the firing rates of transitions with respect to the global the marking of the net system.

3.1. Definition of Π^2-nets

In order to introduce the additional requirement for a Π^2-net to be a Π-net, we illustrate on the Π-net of Fig. 1 how one characterizes by linear combinations of places the virtual client states. Let us focus on the batch jobs. A job is initially idle. When firing t_1 it enters a computing stage followed by a printing stage initiated by firing t_2 and terminates by firing t_3. Thus a job has three states: *idle, computing* and *printing*. The characterization of the states *computing* and *printing* is easy. For instance p_2 is a witness of the number of computing jobs: indeed the start of a job (t_1) puts a token in this place, the beginning of the printing stage (t_2) consumes a token in this place and the firing of any other transition does not modify

the marking of p_2. Similarly p_3 is a witness of the state *printing*. But the characterization of the *idle* state is more intricate: for instance p_1 is not a witness of this state as its marking is modified by t_4 related to another change of state. Nevertheless the linear combination $-(p_2 + p_3)$ is a witness of the *idle* state: indeed given a marking **m**, the end of a job (t_3) increases by one unit $-(\mathbf{m}[p_2] + \mathbf{m}[p_3])$, the start of a job ($t_2$) decreases by one unit $-(\mathbf{m}[p_2] + \mathbf{m}[p_3])$ and the firing of any other transition does not modify $-(\mathbf{m}[p_2] + \mathbf{m}[p_3])$. In order to formally define these linear combinations of places, we generalize the notion of flow. We define a *partial flow* as a pair (**s**, **g**) of $\mathbb{Q}$ vectors on places (**s**, the "solution") and transitions (**g**, the "constraints"), such that $\mathbf{s} \cdot \mathbf{C} = \mathbf{g}$. This is effectively a generalization of flows, for which **g** is always **0**. Thanks to this definition, we characterize the clients with a partial flow where the constraint vector **g** is a vector such that $\mathbf{g}[t] = 1$ if t adds a client to the "state", $\mathbf{g}[t] = -1$ if t removes a client and $\mathbf{g}[t] = 0$ otherwise. The Π^2-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. For the rest of the paper we set $^\bullet\mathbf{r} = \{t \in \mathcal{T} | \mathbf{o}(t) = \mathbf{r}\}$ and $\mathbf{r}^\bullet = \{t \in \mathcal{T} | \mathbf{i}(t) = \mathbf{r}\}$ for every $\mathbf{r} \in \mathcal{R}(\mathcal{T})$.

Definition 7 (Π^2-net). A Π^2-net (parametric Π-net) is a Π-net such that for every $\mathbf{r} \in \mathcal{R}(\mathcal{T})$, there is $\mathbf{a_r} \in \mathbb{Q}^{n_p}$ such that

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

where **C** is the incidence matrix of the net (note that this excludes transitions t with $\mathbf{i}(t) = \mathbf{o}(t)$).

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

The Π-net of Fig. 2 is not a Π^2-net. To see that, let us set $\mathbf{r_1} = \{p_1\}$, therefore $^\bullet\mathbf{r_1} = \{t_4\}$ 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 Π^2-net.

Introducing the stochastic version of Π^2-nets, the explicit relation on states of virtual clients in Π^2-nets, allows us to go one step further: we can express the dependency of the firing rate of a transition t_j with respect to the global state of the components different from the one of t_j. This kind of dependency, introduced by functions $\rho_{\mathcal{C}(t)}$ in the definition below, cannot be taken into account in the framework of Π-nets.

Definition 8 ($S\Pi^2$-net). A Stochastic Π^2-net ($S\Pi^2$-net) is a Π^2-net such that the firing rate of a transition t in the marking **m** is given by

$$\mu(t, \mathbf{m}) = \mu(\mathbf{i}(t)) \cdot \rho_{\mathcal{C}(t)}((\mathbf{a_{r''}} \cdot \mathbf{m})_{\mathbf{r''} \notin \mathcal{C}(t)}) \cdot \frac{\psi(\mathbf{m} - \mathbf{i}(t))}{\phi(\mathbf{m})} \mathbf{P}[\mathbf{i}(t), \mathbf{o}(t)]. \tag{8}$$

Positive, real valued functions $\rho_{\mathcal{C}(t)}((\mathbf{a_{r''}} \cdot \mathbf{m})_{\mathbf{r''} \notin \mathcal{C}(t)})$ make possible a homogeneous dependency of the transitions of the component $\mathcal{C}(t)$ with respect to the state of the virtual clients in the other components, represented by the sequence $(\mathbf{a_{r''}} \cdot \mathbf{m})_{\mathbf{r''} \notin \mathcal{C}(t)}$.

The Π-net of Fig. 1 is a $S\Pi^2$-net. We have six input vectors $\mathbf{r}$, belonging to two classes (we use the same notation $\mathcal{C}(t)$ for the classes of transitions and the classes of their input/output vectors): $\mathcal{C}_1 = \{\mathbf{r_1} = [1, 0, 0, 0, 0, 0], \mathbf{r_2} = [0, 1, 0, 0, 0, 0], \mathbf{r_3} = [0, 0, 1, 0, 0, 0]\}$, $\mathcal{C}_2 = \{\mathbf{r_4} = [1, 0, 0, 1, 0, 0],$ $\mathbf{r_5} = [0, 0, 0, 0, 1, 0], \mathbf{r_6} = [0, 0, 0, 0, 0, 1]\}$. The $\mathbf{a_r}$ vectors are

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

$$\mathbf{a_{r_4}} = [0, 0, 0, 1, 0, 0], \qquad \mathbf{a_{r_5}} = [0, 0, 0, 0, 1, 0], \qquad \mathbf{a_{r_6}} = [0, 0, 0, 0, 0, 1].$$

Let us assume that the rates of t_1, t_2 and t_3 depend on the load of t_5 (and t_6) in such a way that if the marking of p_5 is greater than K_5, t_1, t_2, t_3 cannot fire. Moreover, suppose that the rates of t_1, t_2, t_3 decrease linearly from μ_M to μ_m with the marking of p_5 varying from 0 to K_5. Here we want to model some weak priority of interactive tasks over batch jobs. We emphasize that the $S\Pi$-net model does not allow such a modelling. With $S\Pi^2$-nets, the above dependency is straightforwardly defined:

$$\rho_{\mathcal{C}(t_1)}((\mathbf{a_{r''}} \cdot \mathbf{m})_{\mathbf{r''} \notin \mathcal{C}(t_1)}) = \begin{cases} 0, & \text{if } \mathbf{m}[p_5] \geq K_5, \\ \dfrac{\mu_m - \mu_M}{K_5} \cdot \mathbf{m}[p_5] + \mu_M, & \text{if } 0 \leq \mathbf{m}[p_5] < K_5. \end{cases}$$

We have still a PF steady-state distribution since $\mathbf{a_{r_5}} \cdot \mathbf{m} = \mathbf{m}[p_5]$.

3.2. A product-form for $S\Pi^2$-nets

We establish the first result about $S\Pi^2$-net, which states that a $S\Pi^2$-net is a rate-insensitive PF-$S\Pi$-net.

Theorem 9. *For any transition rates, the steady-state distribution of a $S\Pi^2$-net has the product-form*

$$\pi(\mathbf{m}) = \frac{1}{G} \cdot \phi(\mathbf{m}) \cdot \prod_{\mathbf{r} \in \mathcal{R}(\mathcal{T})} \left(\frac{v(\mathbf{r})}{\mu(\mathbf{r})} \right)^{\mathbf{a_r} \cdot \mathbf{m}}, \quad \forall \mathbf{m} \in RS(\mathbf{m}_0), \tag{9}$$

where G is a normalization constant and v is a solution of Eq. (3).

Proof. The steady-state distribution π of an ergodic CTMC with state space S and generator $\mathbf{Q}$ satisfies the so-called *Cut Balance Equations* (CBE) [29]

$$\sum_{s \in U} \pi(s) \sum_{s' \in \bar{U}} q(s, s') = \sum_{s' \in \bar{U}} \pi(s') \sum_{s \in U} q(s', s) \tag{10}$$

for any subset U and $\bar{U} = S \setminus U$ of S. A classical method [37] to prove that π may have a PF-solution is to group (10) with respect to some partition of S and to find a PF-solution which satisfies this other set of equations, usually termed as *Local Balance Equations* (LBEs). Obviously, if some π satisfies the LBE, then it also satisfies the CBE (but the converse is usually false). We follow this approach here, and we start from the so-called Group Local Balance Equation (GLBE) (11) corresponding to group the CBE with respect to a given vector $\mathbf{r}$ and a marking $\mathbf{m}$

$$\pi(\mathbf{m}) \sum_{t \in \mathbf{r}^\bullet} \mathbf{Q}[\mathbf{m}, \mathbf{m} - \mathbf{i}(t) + \mathbf{o}(t)] = \sum_{t' \in {}^\bullet \mathbf{r}} \pi(\mathbf{m} + \mathbf{i}(t') - \mathbf{o}(t')) \mathbf{Q}[\mathbf{m} + \mathbf{i}(t') - \mathbf{o}(t'), \mathbf{m}]. \tag{11}$$

For simplicity, we first establish the proof with all functions $\rho_{\mathcal{C}(t)}((\mathbf{a_{r''}} \cdot \mathbf{m})_{\mathbf{r''} \notin \mathcal{C}(t)})$ equal to 1.

From the structure of the rates $\mathbf{Q}[\mathbf{m}, \mathbf{m}']$, we derive the equivalent relation

$$\pi(\mathbf{m}) \sum_{t \in \mathbf{r}^\bullet} \mu(\mathbf{i}(t)) \frac{\psi(\mathbf{m} - \mathbf{i}(t))}{\phi(\mathbf{m})} \mathbf{P}[\mathbf{i}(t), \mathbf{o}(t)]$$

$$= \sum_{t' \in {}^\bullet\mathbf{r}} \pi(\mathbf{m} + \mathbf{i}(t') - \mathbf{o}(t'))\mu(\mathbf{i}(t')) \frac{\psi(\mathbf{m} - \mathbf{o}(t'))}{\phi(\mathbf{m} + \mathbf{i}(t') - \mathbf{o}(t'))} \mathbf{P}[\mathbf{i}(t'), \mathbf{o}(t')]$$

and, since $\mathbf{o}(t') = \mathbf{i}(t) = \mathbf{r}$,

$$\pi(\mathbf{m}) \sum_{t \in \mathbf{r}^\bullet} \mu(\mathbf{r}) \frac{\psi(\mathbf{m} - \mathbf{r})}{\phi(\mathbf{m})} \mathbf{P}[\mathbf{r}, \mathbf{o}(t)] = \sum_{t' \in {}^\bullet\mathbf{r}} \pi(\mathbf{m} + \mathbf{i}(t') - \mathbf{r})\mu(\mathbf{i}(t')) \frac{\psi(\mathbf{m} - \mathbf{r})}{\phi(\mathbf{m} + \mathbf{i}(t') - \mathbf{r})} \mathbf{P}[\mathbf{i}(t'), \mathbf{r}],$$

but $\sum_{t \in \mathbf{r}^\bullet} \mathbf{P}[\mathbf{r}, \mathbf{o}(t)] = 1$, so that (11) is equivalent to

$$\frac{\pi(\mathbf{m})\mu(\mathbf{r})}{\phi(\mathbf{m})} = \sum_{t' \in {}^\bullet\mathbf{r}} \pi(\mathbf{m} + \mathbf{i}(t') - \mathbf{r}) \frac{\mu(\mathbf{i}(t'))}{\phi(\mathbf{m} + \mathbf{i}(t') - \mathbf{r})} \mathbf{P}[\mathbf{i}(t'), \mathbf{r}]. \tag{12}$$

Let us prove that the expression of $\pi(\mathbf{m})$ given in the theorem verifies this last equation. To this aim, we introduce in (12) the proposed expression, and after simplification with $1/G \prod_{\mathbf{r}' \in \mathcal{R}(\mathcal{T})} [v(\mathbf{r}')/\mu(\mathbf{r}')]^{\mathbf{a}_{\mathbf{r}'} \cdot \mathbf{m}}$, we obtain the equivalent equation

$$\mu(\mathbf{r}) = \sum_{t' \in {}^\bullet\mathbf{r}} \prod_{\mathbf{r}' \in \mathcal{R}(\mathcal{T})} \left[\frac{v(\mathbf{r}')}{\mu(\mathbf{r}')} \right]^{\mathbf{a}_{\mathbf{r}'} \cdot (\mathbf{i}(t') - \mathbf{o}(t'))} \mu(\mathbf{i}(t'))\mathbf{P}[\mathbf{i}(t'), \mathbf{r}]. \tag{13}$$

From $\mathbf{i}(t') - \mathbf{o}(t') = -\mathbf{C}[\mathcal{P}, t']$ and the definition of $\mathbf{a}_{\mathbf{r}}$, only two terms ($\mathbf{r}' = \mathbf{i}(t')$ and $\mathbf{r}' = \mathbf{r}$) are different from 1 in each product, so we get

$$\mu(\mathbf{r}) = \sum_{t' \in {}^\bullet\mathbf{r}} \frac{\mu(\mathbf{r})}{v(\mathbf{r})} \frac{v(\mathbf{i}(t'))}{\mu(\mathbf{i}(t'))} \mu(\mathbf{i}(t'))\mathbf{P}[\mathbf{i}(t'), \mathbf{r}] \tag{14}$$

and, after rewriting, we obtain

$$v(\mathbf{r}) = \sum_{t' \in {}^\bullet\mathbf{r}} v(\mathbf{i}(t'))\mathbf{P}[\mathbf{i}(t'), \mathbf{r}], \tag{15}$$

which are exactly the traffic equations of the net since if $t' \notin {}^\bullet\mathbf{r}$, $\mathbf{P}[\mathbf{i}(t'), \mathbf{r}] = 0$.

We now extend the proof to general ρ functions. Starting again from Eq. (14) modified with functions ρ, we have

$$\mu(\mathbf{r})\rho_{\mathcal{C}(t)}((\mathbf{a}_{\mathbf{r}''} \cdot \mathbf{m})_{\mathbf{r}'' \notin \mathcal{C}(t)}) = \sum_{t' \in {}^\bullet\mathbf{r}} \frac{\mu(\mathbf{r})}{v(\mathbf{r})} \frac{v(\mathbf{i}(t'))}{\mu(\mathbf{i}(t'))} \mu(\mathbf{i}(t'))\rho_{\mathcal{C}(t')}((\mathbf{a}_{\mathbf{r}''} \cdot (\mathbf{m} + \mathbf{i}(t') - \mathbf{r}))_{\mathbf{r}'' \notin \mathcal{C}(t')})\mathbf{P}[\mathbf{i}(t'), \mathbf{r}].$$

$$\tag{16}$$

In fact, if $t' \in {}^\bullet\mathbf{r}$ and $\mathbf{r}'' \notin \mathcal{C}(t')$, $\mathbf{a}_{\mathbf{r}''} \cdot (\mathbf{m} + \mathbf{i}(t') - \mathbf{r}) = \mathbf{a}_{\mathbf{r}''} \cdot \mathbf{m}$. Indeed, $\mathcal{C}(t') = \mathcal{C}(\mathbf{r})$ and $\mathbf{r}'' \neq \mathbf{r}$ for all $\mathbf{r}''$. Hence, $\mathbf{r}'' \neq \mathbf{i}(t')$ and $\mathbf{r}'' \neq \mathbf{o}(t')$ and $\mathbf{a}_{\mathbf{r}''} \cdot (\mathbf{m} + \mathbf{i}(t') - \mathbf{r}) = \mathbf{a}_{\mathbf{r}''} \cdot \mathbf{m} - \mathbf{a}_{\mathbf{r}''} \cdot \mathbf{C}[\mathcal{P}, t'] = \mathbf{a}_{\mathbf{r}''} \cdot \mathbf{m}$ from the definition of the $\mathbf{a}_{\mathbf{r}}$ vectors. We may now simplify the two terms of Eq. (16) with $\rho_{\mathcal{C}(t)}((\mathbf{a}_{\mathbf{r}''} \cdot \mathbf{m})_{\mathbf{r}'' \notin \mathcal{C}(t)})$ because $\mathcal{C}(t) = \mathcal{C}(t')$ for all $t' \in {}^\bullet\mathbf{r}$, and we obtain Eq. (15). $\qquad \square$

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

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

3.3. Characterization of the rate-insensitive PF-SΠ-nets

GLBE (11) is an essential ingredient to find PF-solution for SPN, and the $S\Pi$-nets for which a PF-solution exists, verify the GLBE. The next theorem proves that a rate-insensitive PF-$S\Pi$-net is a $S\Pi^2$-net. Gathering Theorems 9 and 10, we thus establish that the $S\Pi^2$-nets are exactly the rate-insensitive PF-$S\Pi$-nets.

Theorem 10. *Let* $(\mathcal{P}, \mathcal{T}, W, \mu, \mathbf{m_0})$ *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*

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

satisfies the group local balance equation (11) *for any* $(\mu(\mathbf{r}))_{\mathbf{r}\in\mathcal{R}(\mathcal{T})}$, *then we have*

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

Proof. For simplicity, and without loss of generality (see the proof of Theorem 9), we may assume that $\phi(\mathbf{m}) = \rho_{\mathcal{C}(t)}((\mathbf{a_{r''}}\cdot\mathbf{m})_{\mathbf{r''}\notin\mathcal{C}(t)}) = 1$ for any $\mathbf{m}$ and t.

The GLBE for a given $\mathbf{m}$ with respect to a given $\mathbf{r}$ are (see (13))

$$\mu(\mathbf{r}) = \sum_{t\in{}^{\bullet}\mathbf{r}}\prod_{\mathbf{r'}\in\mathcal{R}(\mathcal{T})}\left[\frac{v(\mathbf{r'})}{\mu(\mathbf{r'})}\right]^{-\mathbf{a_{r'}}\cdot\mathbf{C}[\mathcal{P},t]}\mu(\mathbf{i}(t))\mathbf{P}[\mathbf{i}(t), \mathbf{r}], \tag{17}$$

since $a_{\mathbf{r'}}\cdot(\mathbf{i}(t) - \mathbf{o}(t)) = -\mathbf{a_{r'}}\cdot\mathbf{C}[\mathcal{P}, t]$.

Let $b_t = \prod_{\mathbf{r'}\in\mathcal{R}(\mathcal{T})}[v(\mathbf{r'})]^{-\mathbf{a_{r'}}\cdot\mathbf{C}[\mathcal{P},t]}\mathbf{P}[\mathbf{i}(t), \mathbf{r}]$. $\forall t, b(t) > 0$ since v is a solution of the traffic equations. We define the vectors $\gamma(t)$ and γ_0 in the following way:

$$\gamma(t)[\mathbf{r'}] = \begin{cases} a_{\mathbf{r'}}\cdot\mathbf{C}[\mathcal{P}, t] & \text{if } \mathbf{r'}\neq\mathbf{i}(t), \\ a_{\mathbf{r'}}\cdot\mathbf{C}[\mathcal{P}, t] + 1 & \text{if } \mathbf{r'} = \mathbf{i}(t) \end{cases} \quad \text{and} \quad \gamma_0[\mathbf{r'}] = \begin{cases} 1 & \text{if } \mathbf{r'} = \mathbf{r}, \\ 0 & \text{otherwise.} \end{cases}$$

Eq. (17) is then equivalent to

$$\prod_{\mathbf{r'}\in\mathcal{R}(\mathcal{T})}[\mu(\mathbf{r'})]^{\gamma_0[\mathbf{r'}]} - \sum_{t\in{}^{\bullet}\mathbf{r}}b_t\cdot\prod_{\mathbf{r'}\in\mathcal{R}(\mathcal{T})}[\mu(\mathbf{r'})]^{\gamma(t)[\mathbf{r'}]} = 0. \tag{18}$$

Grouping terms by identical γ, we have

$$\left(1 - \sum_{\substack{t\in{}^\bullet\mathbf{r}\\\gamma(t)=\gamma_0}} b_t\right)\prod_{\mathbf{r}'\in\mathcal{R}(\mathcal{T})}[\mu(\mathbf{r}')]^{\gamma_0[\mathbf{r}']} - \sum_{\gamma\neq\gamma_0}\left(\sum_{\substack{t\in{}^\bullet\mathbf{r}\\\gamma(t)=\gamma}} b_t\right)\prod_{\mathbf{r}'\in\mathcal{R}(\mathcal{T})}[\mu(\mathbf{r}')]^{\gamma(t)[\mathbf{r}']} = 0.$$

It is well-known that if $\mathcal{A}$ is a finite set of *different* vectors $\alpha = (\alpha_1, \ldots, \alpha_n)$ of $\mathbb{R}^n$, and $(a_\alpha)_{\alpha\in\mathcal{A}}$ a family of reals, then $\forall(X_1, \ldots, X_n) \in (\mathbb{R}^+)^n$, $\sum_{\alpha\in\mathcal{A}} a_\alpha X_1^{\alpha_1}, \ldots, X_n^{\alpha_n} = 0 \Rightarrow \forall\alpha \in \mathcal{A}, a_\alpha = 0$. Applying this result with $\mathcal{A} = \{\gamma(t)|t \in \mathcal{T}\}\bigcup\{\gamma_0\}$, $n = |\mathcal{R}(\mathcal{T})|$,

$$a_\alpha = \left(1 - \sum_{\substack{t\in{}^\bullet\mathbf{r}\\\gamma(t)=\gamma_0}} b_t\right)$$

if $\alpha = \gamma_0$ and

$$a_\alpha = \left(\sum_{\substack{t\in{}^\bullet\mathbf{r}\\\gamma(t)=\gamma}} b_t\right)$$

if $\alpha = \gamma(t)$, for $X_{\mathbf{r}'} = \mu(\mathbf{r}')$, we get

$$\sum_{\substack{t\in{}^\bullet\mathbf{r}\\\gamma(t)=\gamma}} b_t = 0.$$

But, since $b_t > 0$ for all t, the set $\{t \in {}^\bullet\mathbf{r}|\gamma(t) = \gamma \neq \gamma_0\}$ is empty, so that $\forall t \in {}^\bullet\mathbf{r}$, $\gamma(t) = \gamma_0$. We can now evaluate the $a_{\mathbf{r}'} \cdot \mathbf{C}[\mathcal{P}, t]$ numbers (let us remind that $t \in {}^\bullet\mathbf{r}$ means $\mathbf{o}(t) = \mathbf{r}$):

if $\mathbf{i}' \neq \mathbf{r}(t)$ and $\mathbf{r}' \neq \mathbf{o}(t)$, $\quad a_{\mathbf{r}'} \cdot \mathbf{C}[\mathcal{P}, t] = \gamma(t)[\mathbf{r}'] = 0$;

if $\mathbf{r}' = \mathbf{i}(t)$ (hence $\mathbf{r}' \neq \mathbf{o}(t)$), $\quad a_{\mathbf{r}'} \cdot \mathbf{C}[\mathcal{P}, t] + 1 = \gamma(t)[\mathbf{r}'] = 0$;

if $\mathbf{r}' \neq \mathbf{i}(t)$ and $\mathbf{r}' = \mathbf{o}(t)$, $\quad a_{\mathbf{r}'} \cdot \mathbf{C}[\mathcal{P}, t] = \gamma(t)[\mathbf{r}'] = 1$.

which concludes the proof. $\qquad\square$

To be complete, we have established in [35] a connection between $S\Pi^2$-nets and Product Process of K Competing Markov Chains MC_k (PPCMC) introduced by Boucherie in [5]. Boucherie has studied nets with a product-form solution which are not $S\Pi$-nets. He has showed on different case studies their

transformation into semantically equivalent $S\Pi$-nets (i.e. generating the same CTMC). We have showed that these equivalent $S\Pi$-nets are $S\Pi^2$-nets.

3.4. Comparison with the previous characterization

In [13], Coleman et al. proposed another characterization of the rate-insensitive PF-$S\Pi$-nets expressed in terms of batch queuing networks which are equivalent to $S\Pi$-nets. We first recall in extenso their characterization:

> "Suppose that the initial state $\mathbf{n}_0$ is large enough so that the set of possible firing vectors resulting from transition sequences that take the state of the network from $\mathbf{n}_0$ back to itself is the set of *T*-invariants. Then the equilibrium distribution factorizes into a product-form over the nodes for all values $\Xi(\mathbf{a})$ if and only if there is a one-to-one correspondence between states of the network and count vectors."

In this theorem, the family of values $\Xi(\mathbf{a})$ corresponds the family $\mu(\mathbf{i})$ presented here and the count vectors are exactly what we call the virtual client states. We now compare the two characterizations:

- The previous characterization points out that, given a state of the net, there is a single set of current virtual client states corresponding to it. Our characterization shows that furthermore this correspondence is *necessarily* given by a linear application.
- The previous characterization requires that the initial marking must be sufficiently large whereas our characterization is valid for any initial marking.
- The previous characterization must be checked at the reachability graph level with a procedure whose time complexity is exponential with respect to the size of the net in the worst case. Proposition 12 shows that this negative result holds even for 1-safe $S\Pi$-nets. Here we provide a polynomial time algorithm to check our characterization.
- This new characterization has enabled us to enlarge the functional dependencies of the rates of the transitions. Otherwise stated, the $S\Pi^2$-nets family is a strict superset of the $S\Pi$-nets family except for the pathological cases.

4. Qualitative analysis of Π-nets and Π^2-nets

Since the central problems related to PNs (liveness, boundedness, reachability, coverability, etc.) have high complexity lower bounds, the consideration of some net subclasses, enjoying particular properties has quickly appeared as a mean to cope with this complexity. Forty years of research in the PN area, have accumulated a large amount of qualitative results and several classes of PN have been identified, particularly: State Machines (SM), Marked Graphs (MG) (also known as T-Systems (TS)) and Extended Free-Choice (EFC) nets. It is consequently a standard approach to situate a new class of PNs with respect to these "classical" subclasses: this may allow us to easily derive some of their properties (boundedness, deadlock freeness, etc.). In fact, it was shown (see [43] for instance) that Π-nets cannot be directly compared to standard families of PN. We strengthen below these negative results, by complexity results which constitute our third contribution.

ARTICLE IN PRESS

It is clear that if there is an enabled transition in any FR*-class in the initial marking of a Π-net, then the net system is live. This condition is obviously only a sufficient condition. In fact, checking liveness seems not easier for Π-nets, and even for 1-safe Π^2-nets, than for many other classes of Petri nets (a 1-safe marked Petri net is a (bounded) marked net with at most one token in every place of every reachable marking). We have shown in Section 2.2 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 complex problem. Indeed, the next proposition gives some insight into this point. We recall that for general Petri nets, Lipton's result [32] implies a $2^{O(\sqrt{n})}$ lower bound space complexity (with n the size of the problem, see Section 2.2) for the liveness problem (see [16,15] for recent surveys on decidability problems for Petri nets). In fact, we are able to complement these results for the Π-nets and Π^2-nets classes, although the exact complexity of the reachability/liveness for Π^2-nets still remains an open problem. Moreover, it has been also shown in [11] 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^2-nets.

Proposition 11. *The liveness problem for 1-safe Π^2-nets is NP-hard.*

Proof. To prove it, we reduce in polynomial time the 3SAT problem to the liveness problem for Π^2-nets, following the idea first presented in [26]. The 3SAT problem is a well known NP-complete problem. We have K logical formulae $C_1, \ldots, C_K$, each one being a disjunction of three boolean variables v_i or their negation ($\bar{v}_i$), from a set of I variables: for instance, $C_K = v_1 \vee \bar{v}_3 \vee v_6$. The 3SAT problem can be stated as follows: is there a set of values for $v_1, \ldots, v_I$ such that $C_1 \wedge C_2 \wedge \cdots \wedge C_K$ is true? We explain the reduction through the example $C_1 = v_1 \vee \bar{v}_2 \vee v_3$, $C_2 = v_2 \vee v_3 \vee v_4$ ($K = 2$, $I = 4$) (Fig. 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 $p_{C_k,i}$ (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 $p_{C_k,i}$ and one arc from $p_{C_k i}$ to t_i. In contrast if $\bar{v}_i$ is in C_k (like $\bar{v}_2$ 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_{C_k} 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 Π^2-net and not only a Π-net. We choose an initial uniform affectation *true* for the v_i. So the initial marking is chosen as follows: there is one token in p_i and one token in place $p_{C_k i}$ if v_i occurs in C_k.

We can easily verify that the net system is a 1-safe Π^2-net and that all transitions except possibly t_{s1}, t_{s2} are live. Clearly, the formula is true for some set of boolean values of variables iff the transition t_{s1} is live (and consequently t_{s2} is also live). □

Thus, there is still an open problem for Π^2-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.

Proposition 12.

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.*

Fig. 3. Reduction of 3SAT to liveness in 1-safe Π^2-nets.

Proof.

1. Let us first prove that the reachability problem for 1-safe Π^2-nets (RP for short) is PSPACE complete. Since RP for 1-safe Petri nets is in PSPACE [11], we only prove that RP is PSPACE hard. To this end, we reduce in polynomial time the termination problem of deterministic Türing Machines (DTM) with finite length tape to RP. Let M be a DTM with a tape of length n on an alphabet $A = \{a_1, \ldots, a_m\}$ (including the blank character), a set Q of states (with Q_F the subset of final states), and a (partially defined) transition function $\delta(q, a) = (q', a', d)$ with $d \in \{n, l, r\}$ (no move, left move, right move). A configuration of M is a tuple $\langle \mathbf{a}, q, i \rangle$ with $\mathbf{a} \in A^n$ the current content of the tape, $q \in Q$ the current state and $1 \leq i \leq n$ the current position of the head. A change of configuration is denoted by $\langle \mathbf{a}, q, i \rangle \rightarrow_M \langle \mathbf{a}', q', i' \rangle$.

   We encode M by a Petri net system $\mathcal{S}(M) = (\mathcal{N}, \mathbf{m}_0)$. The set of places is partitioned into three subsets: $\mathcal{P}_{\text{tape}} = \{c_{i,j} | 1 \leq i \leq n, 1 \leq j \leq m\}$ encodes the possible values of the cells $\{i\}_{1 \leq i \leq n}$ of the tape; Q encodes the possible states of M; $\mathcal{P}_{\text{head}} = \{h_i | 1 \leq i \leq n\}$ encodes the possible positions of the head of M. Given q a state of M, i a position of the head and j such that a_j is the possible value of the cell i, we define $t_{q,ij}$ a transition of $\mathcal{N}$ which represents the change of configuration. For instance (see Fig. 4), if $\delta(q, a_j) = (q', a_{j'}, l)$, the transition $t_{q,i,j}$ $(i > 1)$ is defined as follows: there are three input arcs from q, h_i and c_{ij} to $t_{q,i,j}$ and there are three output arcs from $t_{q,i,j}$ to q', h_{i-1} and $c_{i,j'}$. Let us call $\mathbf{m}$ an appropriate marking iff $\sum_{q \in Q} \mathbf{m}(q) = 1$, $\sum_{1 \leq i \leq n} \mathbf{m}(h_i) = 1$ and $\forall 1 \leq i, \leq n \sum_{1 \leq j \leq m} \mathbf{m}(c_{i,j}) = 1$. There is an obvious one-to-one correspondence between the configurations of M and the appropriate markings of $\mathcal{N}$. We choose $\mathbf{m}_0$ as the appropriate marking corresponding to the initial configuration

ARTICLE IN PRESS

Fig. 4. Encoding a DTM with a Petri net.

of M. Given $\mathbf{m}$ an appropriate marking, we note $\langle \mathbf{a_m}, q_{\mathbf{m}}, i_{\mathbf{m}} \rangle$ the configuration corresponding to $\mathbf{m}$. Thus by construction, one has $\mathbf{m} \rightarrow^{t_{q,i,j}} \mathbf{m}'$ iff $\mathbf{m}'$ is an appropriate marking and $q = q_{\mathbf{m}}$ and $i = i_{\mathbf{m}}$ and $\mathbf{a}(i) = a_j$ and $\langle \mathbf{a_m}, q_{\mathbf{m}}, i_{\mathbf{m}} \rangle \rightarrow_M \langle \mathbf{a_{m'}}, q_{\mathbf{m}'}, i_{\mathbf{m}'} \rangle$.

This property shows that $\mathcal{S}(M)$ is an exact simulation of M. So the termination problem for M is reduced to a coverability problem of $\mathcal{S}(M)$, i.e. to find a reachable marking $\mathbf{m}$ s.t. $\mathbf{m}(q) = 1$ for some $q \in Q_F$. To get a reachability problem we may assume, without loss of generality, that M has moreover the following properties: the set Q_F is a singleton $\{q_F\}$; in q_F, M is "reset", i.e. the head of M is on the first case of the tape and $\forall i$, $b_i = a_1$, and there is no transition from q_F. Then the termination of M is now reduced to the reachability of the appropriate marking (denoted $\mathbf{m}_F$) corresponding to this final configuration.

Of course $\mathcal{S}$ is not necessarily 1-safe Π-net. Starting from $\mathcal{S}$, we define a 1-safe Π-net system $\mathcal{S}'$, adding for each transition $t_{q,i,j}$ of $\mathcal{N}$, its "reverse" transition denoted by $t'_{q,i,j}$ (the firings of the "reverse" transition t' "undoes" the effect of the firing of t, i.e., the input places of t are the output places of t' and vice versa). Observe that in $\mathcal{S}'$, starting from an appropriate marking, we reach only appropriate markings.

We claim that the termination problem of M is reduced to the reachability of $\mathbf{m}_F$ in $\mathcal{S}'$. Since any firing sequence of $\mathcal{S}$ is a firing sequence of $\mathcal{S}'$, if M terminates then $\mathbf{m}_F$ is reachable in $\mathcal{S}'$. We now prove that if $\mathbf{m}_F$ is reachable in $\mathcal{S}'$, then $\mathbf{m}_F$ is reachable in $\mathcal{S}$ and so M terminates. Assume that $\mathbf{m}_0 \rightarrow^{\sigma'} \mathbf{m}_F$, with σ' a firing sequence of $\mathcal{N}$. We claim that there is a firing sequence σ of $\mathcal{N}$ s.t. $\mathbf{m}_0 \rightarrow^{\sigma} \mathbf{m}_F$. The proof is done by induction on $|\sigma'| = k$. If $k = 0$, $\mathbf{m}_0 = \mathbf{m}_F$. If $k > 0$ and if there is no "reverse" transition in σ' the $\sigma = \sigma'$ otherwise let us denote by $t'_{q,i,j}$ the last reverse transition of σ', so that $\mathbf{m}_0 \rightarrow^{ut'_{q,i,j}} \mathbf{m} \rightarrow^{v} \mathbf{m}_F$ and v is a sequence of $\mathcal{N}$. Then we have first $|v| \geq 1$, because if v is empty then $q = q_F$ the definition of $t_{q,i,j}$ implies that $\delta(q_F, a_j)$ is defined which contradicts our hypotheses on M. So $v = t_{q',i'j'} v'$. Consequently, we have $\mathbf{m} \rightarrow^{t_{q,i,j}}$ and $\mathbf{m} \rightarrow^{t_{q',i',j'}}$. Since $\mathbf{m}$ is an appropriate marking the above property and the fact that M is deterministic ensure that $\langle q, i, j \rangle = \langle q', i', j' \rangle$. So, we have $\mathbf{m}_0 \rightarrow^{ut'_{q,i,j} t_{q,i,j} v'} \mathbf{m}_F$, implying $\mathbf{m}_0 \rightarrow^{uv'} \mathbf{m}_F$ with $|uv'| < k$ and uv' satisfies the induction hypothesis.

The reduction to the liveness problem may be done in the same way, extending the translated Petri net system $\mathcal{S}'$ with extra places and transitions.

2. For symmetric nets systems, we know [10,34] that the reachability problem is EXPSPACE complete. A net is symmetric iff for every transition t, the reverse t' of t is a transition of the net. 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 set, because from any closed T-semiflow of transitions to which t belongs, one can pick the other transitions of this flow and build a sequence (following the input/output bags) such that this sequence is firable iff the reverse transition is firable with the same reached marking. Thus, the reachability problem for Π-nets is reducible to the one for symmetric nets, hence in EXPSPACE, and finally EXPSPACE complete. □

5. Conclusion

In this paper, we have characterized the class of rate-insensitive PF-Π-nets giving a definitive answer to a question partially solved in [13]. This characterization has two important features. It relies on purely structural conditions (i.e., it is not defined in terms of the reachability graph) and it can be checked in polynomial time with respect to the size of the net. Furthermore this structural characterization has allowed us to extend the previous model with new dependencies between components of the net, thus covering a broader range of applications. We have also studied the complexity of the reachability and liveness problems in Π-nets and Π^2-nets. We have established lower and upper bounds for the complexity of different problems obtaining for some of them an exact characterization. This work has different perspectives that we describe below:

- There are still some open problems about the exact characterization of checking some properties (e.g., the reachability problem in 1-safe Π^2-nets).
- The efficient computation of the normalization constant is closely related to the research of a compact representation of the reachability set of the Petri net. When looking about our complexity results seems that, in the general case, this objective is unrealistic. So we need to identify additional structural properties that must enjoy a Π^2-net in order to obtain such a compact representation.
- Another direction which would enlarge the application of product-form methods would be to use it as part of an approximate method. For instance, given a stochastic Petri net, we could compute the T-semiflows, then we could define T-components of the net with respect to these semi-flows and we should apply the product-form solution to a Π^2-net obtained by a slight transformation of the original net. Obviously, this direction must be investigated in order to examine for which kinds of nets such a method would give accurate (or acceptable) results.
- The implementation of the method is planned in the tool GreatSPN [12]. This will complement the numerous functionalities of this software.
- Whereas this product-form strongly relies on the algebraic nature of Petri net (i.e., T-semiflows and partial flows), we still believe that we can obtain similar but different product-forms for other models like stochastic process algebra [23], queueing networks with resources [27,42] and networks of stochastic automata [39].

References

[1] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, G. Franceschinis, Modelling with Generalized Stochastic Petri Nets, Wiley, New York, 1995.

[2] F. Baskett, K.M. Chandy, R.R. Muntz, F. Palacios, Open, closed and mixed networks of queues with different classes of customers, J. ACM 22 (2) (1975) 248–260.

[3] B. Baynat, Y. Dallery, Approximate analysis of multi-class synchronized closed queueing networks, in: Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MAS-COTS'95), Durham, NC, January 1995.

[4] G. Bolch, S. Greiner, H. de Meer, K.S. Trivedi, Queueing Networks and Markov Chains, Wiley, New York, 1998.

[5] R.J. Boucherie, A characterisation of independence for competing Markov chains with application to stochastic Petri nets, in: Proceedings of the 5th International Workshop on Petri Nets and Performance Models, Toulouse, France, October 19–22, IEEE Computer Society Press, Silverspring, MD, 1993, pp. 117–126.

[6] R.J. Boucherie, M. Sereno, On closed support T-invariants and traffic equations, J. Appl. Prob. (35) (1998) 473–481.

[7] J.P. Buzen, Computational algorithms for closed queueing networks with exponential servers, Commun. ACM 16 (9) (1973) 527–531.

[8] J. Campos, J.M. Colom, H. Jungnitz, M. Silva, Approximate throughput computation of stochastic marked graphs, IEEE Trans. Softw. Eng. 20 (7) (1994) 526–535.

[9] J. Campos, M. Silva, Structural techniques and performance bounds of stochastic Petri nets models. Desing methods based on nets (DEMON), in: G. Rozenberg (Ed.), Advances in Petri Nets, Lecture Notes in Computer Science, vol. 609, Springer, Berlin, 1992, pp. 352–391.

[10] E. Cardoza, R.J. Lipton, A.R. Meyer, Exponential space complete problems for Petri nets and commutative semigroups, in: Proceedings of the 8th Annual Symposium on Theory of Computing, 1976, pp. 50–54.

[11] A. Cheng, J. Esparza, J. Palberg, Complexity results for 1-safe nets, in: Proceedings of the 13th Conference on Foundations of Software Technology and Theoretical Computer Science, Bombay, India, 1993.

[12] G. Chiola, G. Franceschinis, R. Gaeta, M. Ribaudo, GreatSPN1.7: GRaphical Editor and Analyzer for Timed and Stochastic Petri Nets, Perform. Eval. 24 (1–2) (1995) 47–68.

[13] J.L. Coleman, W. Henderson, C.E.M. Pearce, P.G. Taylor, A correspondence between product form batch movement queueing networks and single movement networks, J. Appl. Prob. 34 (1) (1997) 160–175.

[14] J.L. Coleman, W. Henderson, P.G. Taylor, Product form equilibrium distributions and a convolution algorithm for stochastic Petri nets, Perform. Eval. 26 (3) (1996) 159–180.

[15] J. Esparza, Decidability and Complexity of Petri Nets Problems—An Introduction, Springer, Berlin, 1998, pp. 374–428.

[16] J. Esparza, M. Nielsen, Decidability issues for Petri nets—a survey, J. Inform. Process. Cybernet. 30 (3) (1994) 143–160 (former version in Bull. EATCS 52 (1994) 245–262).

[17] D. Frosch, Product form solution for closed synchronized systems of stochastic sequential processes, in: Proceedings of the International Computer Symposium, Taichung, Taiwan, December 1992, pp. 392–402.

[18] W.J. Gordon, G.F. Newell, Closed queueing systems with exponential servers, Oper. Res. 15 (1967) 254–265.

[19] W. Henderson, D. Lucic, P.G. Taylor, A net level performance analysis of stochastic Petri nets, J. Aust. Math. Soc., Ser. B 31 (1989) 176–187.

[20] W. Henderson, C.E.M. Pearce, P.G. Taylor, N.M. van Dijk, Closed queueing networks and batch services, Queueing Syst. 6 (1990) 59–70.

[21] W. Henderson, P.G. Taylor, Aggregation methods in exact performance analysis of stochastic Petri nets, in: Proceedings of the 3rd International Workshop on Petri Nets and Performance Models, Kyoto, Japan, December, IEEE-CS Press, Silverspring, MD, 1989, pp. 12–18.

[22] W. Henderson, P.G. Taylor, Embedded processes in stochastic Petri nets, IEEE Trans. Softw. Eng. 17 (1991) 108–116.

[23] J. Hillston, A Compositional Approach to Performance Modelling, Cambridge University Press, Cambridge, 1996.

[24] J.R. Jackson, Jobshop-like queueing systems, Manage. Sci. 10 (1) (1963) 131–142.

[25] J.R. Jackson, Networks of waiting lines, Oper. Res. 3 (1957) 518–521.

[26] N.D. Jones, L.H. Landweber, Y.E. Lien, Complexity of some problems in Petri nets, Theoret. Comput. Sci. 4 (1977) 277–299.

[27] T.W. Keller, Computer systems models with passive resources, Ph.D. Dissertation, University of Texas, Austin, USA, 1976.

[28] F.P. Kelly, Markov processes and Markov random fields, Bull. Inst. Int. Statist. 46 (1975) 397–404.

[29] F.P. Kelly, Reversibility and Stochastic Networks, Wiley, London, 1979.

[30] A.A. Lazar, T.G. Robertazzi, Markovian Petri net protocols with product form solution, in: Proceedings of the INFO-COM'87, San Francisco, CA, USA, 1987, pp. 1054–1062.

[31] M. Li, N.D. Georganas, Parametric analysis of stochastic Petri nets, in: Proceedings of the Fifth International Conference on Modelling and Tools for Computer Performance Evaluation, Torino, Italy, February, 1991.

[32] R.J. Lipton, The reachability problem requires exponential space, Technical Report, Department of Computer Science, Yale University, 1976.

[33] J. Martinez, M. Silva, A simple and fast algorithm to obtain all invariants of a generalized Petri net, in: Proceedings of the 2nd European Workshop on Application and Theory of Petri Nets, Bad Honnef, West Germany, Springer, Berlin, 1981.

[34] E.W. Mayr, A.R. Meyer, The complexity of the word problems for commutative semigroups and polynomial ideals, Adv. Math. (46) (1982) 305–329.

[35] P. Moreaux, Boucherie product process of competing Markov chains and $S\Pi^2$ nets, May 30, 2002. Internal note available at http://lamsade.dauphine.fr/~moreaux/publications/ppcmcspi2n-tr02.pdf.

[36] T. Murata, Petri nets: properties, analysis, and applications, Proc. IEEE 77 (4) (1989) 541–580.

[37] R.D. Nelson, The mathematics of product form queuing networks, ACM Comput. Surv. 25 (3) (1993) 339–369.

[38] J.L. Peterson, Petri Net Theory and the Modeling of Systems, Prentice-Hall, Englewood Cliffs, NJ, 1981.

[39] B. Plateau, K. Atif, Stochastic automata network for modeling parallel systems, IEEE Trans. Softw. Eng. 17 (10) (1991) 1093–1108.

[40] M. Reiser, S.S. Lavenberg, Mean value analysis of closed multichain queueing networks, J. ACM 27 (2) (1980) 313–322.

[41] W. Reisig, G. Rozenberg (Eds.), Advances in Petri Nets, Lectures on Petri Nets I: Basic models, Lecture Notes in Computer Science, vol. 1491, Springer, Berlin, 1998.

[42] H. Sauer, E.A. MacNair, J.F. Kurose, The research queueing package: past, present, and future, in: Proceedings of the National Computer Conference, AFIPS, 1982.

[43] M. Sereno, Product form and Petri nets, Modelling and analysis of time constrained and hierarchical systems (MATCH), ERB Project CHRX-CT-940452, Performance Models for Discrete Event Systems with Synchronisations: Formalisms and Analysis Techniques, 1998, Chapter 18.

[44] R. Serfozo, Markovian network processes: congestion-dependent routing and processing, Queueing Syst. 5 (1989) 5–36.

[45] M. Silva, Las Redes de Petri en la Automatica y la Informatica, Madrid, Spain, 1985 (in Spanish).

[46] M. Silva, Toward a synchrony theory for P/T nets, in: K. Voos, H.J. Genrich, G. Rozenberg (Eds.), Concurrency and Nets, Springer, Berlin, 1987, pp. 435–460.

Serge Haddad is an old student of the Ecole Normale Supérieure l'Enseignement Technique in mathematics. After passing his PhD degree from University Paris 6 in june 1987 about theory of Petri nets, he became associate professor of this University. Since September 1993, he is a full professor of University Paris-Dauphine. His research interests are focused on the design, the verification and the evaluation of parallel and distributed systems.

Patrice Moreaux is "professeur agrégé de Mathématiques" and received his PhD degree in Computer Science from the University Paris Dauphine in 1996. He is currently associate professor at the Mathematics and Computer Sciences Department of the Université de Reims Champagne-Ardenne, and researcher at Reims and at the LAMSADE laboratory of the Université Paris Dauphine. Its research interests are the modelling, the verification and the performance evaluation in automation, networks, distributed and parallel computer systems.

Matteo Sereno was born in Nocera Inferiore, Italy. He received his Doctor degree in Computer Science from the University of Salerno, Italy, in 1987, and his PhD degree in Computer Science from the University of Torino, Italy, in 1992. He is associate professor at the Computer Science Department of the University of Torino. His current research interests are in the area of performance evaluation of computer systems, modelling of communication networks and parallel architectures, queueing network and stochastic Petri net models.

Manuel Silva received the Industrial-Chemical Engineering degree from the University of Sevilla in 1974 and the postgraduate and PhD degrees in Control Engineering from the Institut National Polytechnique de Grenoble in 1975 and 1978, respectively. In 1978 he started the group of Systems Engineering and Computer Science at the University of Zaragoza, where he was named Professor in Systems Engineering and Automatic Control in 1982. His research interests include modelling, validation, performance evaluation, and implementation of distributed concurrent systems using Petri nets, binary decision graphs, and robots programming and control. He is author of the book Las Redes de Petri en la Automàtica y la Informàtica (AC, 1985), coauthor of the book Practice of Petri Nets in Manufacturing (Chapman & Hall, 1993), and co-editor of one special issue on CIM of the IEEE Trans. on Robotics and Automation. He has been associate editor of the IEEE Trans. on Robotics and Automation and associate editor of the European Journal of Control. He is advisory member of the IEICE Trans. on Fundamentals on Electronics, Communications and Computer Sciences, member of the Steering Committee of the International Conference on Application and Theory of Petri Nets, founder member of the Asociación Española de Robótica, and member of the IEEE, IFAC, and EATCS. He is member of the National Academy of Engineering of Spain.