

Deficiency Zero Petri Nets and Product Form

Jean Mairesse¹ and Hoang-Thach Nguyen¹

LIAFA, Université Paris Diderot - Paris 7, 75205 Paris Cedex 13
 mairesse@liafa.jussieu.fr, hoang-thach.nguyen@liafa.jussieu.fr

Abstract. Consider a Markovian Petri net with race policy. The marking process has a “product form” stationary distribution if the probability of viewing a given marking can be decomposed as the product over places of terms depending only on the local marking. First we observe that the Deficiency Zero Theorem of Feinberg, developed for chemical reaction networks, provides a structural and simple sufficient condition for the existence of a product form. In view of this, we study the classical subclass of free-choice nets. Roughly, we show that the only such Petri nets having a product form are the state machines which can alternatively be viewed as Jackson networks.

1 Introduction

Queueing networks, Petri nets, and chemical reaction networks, are three mathematical models of “networks”, each of them with an identified community of researchers.

In queueing, the existence of “product form” Markovian networks is one of the cornerstones and jewels of the theory. Monographies are dedicated to the subject, e.g. Kelly [17] or Van Dijk [9]. Roughly, the interest lies in the equilibrium behavior of a Markovian queueing network. The existence of such an equilibrium is equivalent to the existence of a stationary distribution π for the queue-length process. In some remarkable cases, π not only exists but has an explicit decomposable shape called “product form”. The interest is two-fold. First, from a quantitative point of view, it makes the explicit computation of π possible, even for large systems. Second, the product form has important qualitative implications, like the “Poisson-Input Poisson-Output” Theorems. Consequently important and lasting efforts have been devoted to the quest for product form queueing networks.

It is attractive and natural to try to develop an analog theory for Markovian Petri nets, with the marking process replacing the queue-length process. There is indeed a continuing string a research on this topic since the 90ies, e.g. [5, 6, 10, 12, 14, 15, 19]. Tools have been developed in the process which build on classical objects of Petri net theory (e.g. closed support T-invariants). The most accomplished results are the ones in [14].

In chemistry and biology, has emerged the model of chemical reaction networks. Such a network is specified by a finite set of reactions between species

of the type “ $2A + B \rightarrow C$ ”, meaning that two molecules of A can interact with one molecule of B to create one molecule of C . The dynamics of such models is either deterministic or stochastic, see [18].

Deterministic models are the most studied ones, they correspond to coupled sets of ordinary differential equations. The most significant result is arguably the Deficiency Zero Theorem of Feinberg [11]. Deficiency Zero is a structural property which can be very easily checked knowing the shape of the reactions in a chemical network. It does not refer to any assumption on the associated dynamics. Feinberg Theorem states that if a network satisfies the Deficiency Zero condition, then the associated deterministic dynamic model has remarkable stability properties. An intermediate result is to prove that a set of non-linear equations (NLE) have a strictly positive solution.

Stochastic models of chemical reaction networks correspond to continuous-time Markov processes of a specific shape. Such models were considered in Chapter 8 of the seminal book by Kelly [17]. There it is proved that if a set of non-linear “traffic equations” (NLTE) have a strictly positive solution then the Markov process has a product form.

How does Feinberg result connect with product form Markovian Petri nets ?

A first observation is that chemical reaction networks and Petri nets are two different descriptions of the same object. This has been identified by different authors in the biochemical community, see for instance [3] and the references therein. Conversely, Petri nets were originally introduced by Carl Adam Petri to model chemical processes, see [22].

A second observation was made recently by Anderson, Craciun, and Kurtz [2]. They observe that the NLE of Feinberg and the NLTE of Kelly are the same. It implies that if a chemical network has deficiency zero, then the stochastic dynamic model has a product form.

In the present paper, we couple the two observations together. The Deficiency Zero condition provides a sufficient condition for a Markovian Petri net to have a product form. The Deficiency Zero condition is equivalent to the criterion known in the Petri net literature [14]. The advantage is that deficiency is easy to compute and handle.

The class of Petri nets whose Markovian version have a product form is an interesting one. It is therefore natural to study how this class intersects with the classical families of Petri nets: state machines and free-choice Petri nets. We use the simplicity of the Deficiency Zero condition to carry out this study.

The central result that we prove is, in a sense, a negative result. We show that within the class of free-choice Petri nets, the only ones which have a product form are closely related to state machines. We also show that the Markovian state machines are “equivalent” to Jackson networks. The latter form the most basic and classical example of product form queueing networks.

A conference version of the present paper appears in [20]. Compared with [20], additional results have been proved: the equivalence between deficiency 0 and the condition of Haddad & al [14] (Prop. 3.9), and the results in Section 3.5.

2 Model

We use the notation $\mathbb{R}^* = \mathbb{R} - \{0\}$. The coordinate-wise ordering of \mathbb{R}^k is denoted by the symbol \leq . We say that $x \in \mathbb{R}^k$ is *strictly positive* if $x_i > 0$ for all i . We denote by $\mathbf{1}_S$ the indicator function of S , that is the mapping taking value 1 inside S and 0 outside.

2.1 Petri nets

Our definition of Petri net is standard, with weights on the arcs.

Definition 2.1 (Petri net). A Petri net is a 6-tuple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, I, O, M_0)$ where:

- $(\mathcal{P}, \mathcal{T}, \mathcal{F})$ is a directed bipartite graph, that is, \mathcal{P} and \mathcal{T} are non-empty and finite disjoint sets, and \mathcal{F} is a subset of $(\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})$;
- $I : \mathcal{T} \rightarrow \mathbb{N}^{\mathcal{P}}$ and $O : \mathcal{T} \rightarrow \mathbb{N}^{\mathcal{P}}$ are such that $[I(t)_p > 0 \Leftrightarrow (p, t) \in \mathcal{F}]$ and $[O(t)_p > 0 \Leftrightarrow (t, p) \in \mathcal{F}]$;
- M_0 belongs to $\mathbb{N}^{\mathcal{P}}$.

The elements of \mathcal{P} are called *places*, those of \mathcal{T} are called *transitions*. The 5-tuple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, I, O)$ is called the *Petri graph*. The vectors $I(t)$ and $O(t)$, $t \in \mathcal{T}$, are called the *input bag* and the *output bag* of the transition t . An element of $\mathbb{N}^{\mathcal{P}}$ is called a *marking*, and M_0 is called the *initial marking*.

Petri nets inherit the usual terminology of graph theory. Graphically, a Petri net is represented by a directed graph in which places are represented by circles and transitions by rectangles. The initial marking is also materialized: if $M_0(p) = k$, then k tokens are drawn inside the circle p . See Figure 1 for an example.

A Petri net is a dynamic object. The Petri graph always remains unchanged, but the marking evolves according to the *firing rule*. A transition t is *enabled* in the marking M if $M \geq I(t)$, then t may *fire* which transforms the marking from M into

$$M' = M - I(t) + O(t).$$

We write $M \xrightarrow{t} M'$. A marking M' is *reachable* from a marking M if there exists a sequence of transitions t_1, \dots, t_k , and a sequence of markings M_1, \dots, M_{k-1} , such that $M \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots \xrightarrow{t_{k-1}} M_{k-1} \xrightarrow{t_k} M'$. We denote by $\mathcal{R}(M)$ the set of markings which are reachable from M .

Definition 2.2 (Marking graph). The marking graph of a Petri net with initial marking M_0 is the directed graph with

- nodes: $\mathcal{R}(M_0)$, arcs: $M \rightarrow M'$ if $\exists t \in \mathcal{T}, M \xrightarrow{t} M'$.

The marking graph defines the state space on which the marking may evolve. Observe that the marking graph may be finite or infinite. In Section 2.3, we will define a *Markovian Petri net* as a continuous-time Markovian process evolving on the marking graph.

The analysis of Petri nets relies heavily on linear algebra techniques, the central object being the incidence matrix.

The *incidence matrix* N of the Petri net $(\mathcal{P}, \mathcal{T}, \mathcal{F}, I, O, M_0)$ is the $(\mathcal{P} \times \mathcal{T})$ -matrix N defined by:

$$N_{s,t} = O(t)_s - I(t)_s. \quad (1)$$

Example. Figure 1 represents a Petri net with places $\{p_1, p_2, p_3, p_4\}$ and transitions $\{t_1, t_2, t_3, t_4\}$. The initial marking is $M_0 = (2, 1, 0, 1)$. The input and output bags are:

$$\begin{aligned} I(t_1) &= (2, 0, 0, 0), O(t_1) = (0, 2, 0, 0), & I(t_2) &= (0, 2, 0, 0), O(t_2) = (2, 0, 0, 0), \\ I(t_3) &= (1, 0, 1, 0), O(t_3) = (0, 1, 0, 1), & I(t_4) &= (0, 1, 0, 1), O(t_4) = (1, 0, 1, 0). \end{aligned}$$

The weights different from 1 are represented on the arcs.

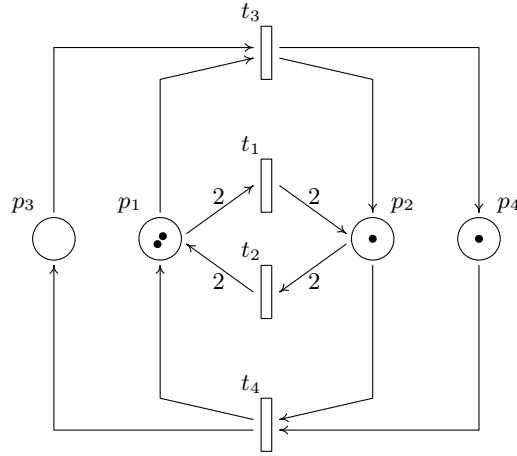


Fig. 1. Petri net.

The reachable markings are M_0 , $M_1 = (0, 3, 0, 1)$, $M_2 = (3, 0, 1, 0)$, and $M_3 = (1, 2, 1, 0)$. The marking graph is represented on Figure 2.

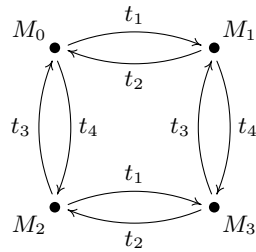


Fig. 2. Marking graph.

The incidence matrix of the Petri net is:

$$N = \begin{pmatrix} -2 & 2 & -1 & 1 \\ 2 & -2 & 1 & -1 \\ 0 & 0 & -1 & 1 \\ 0 & 0 & 1 & -1 \end{pmatrix}. \quad (2)$$

2.2 Deficiency and weak reversibility

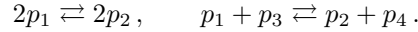
Consider a Petri net $(\mathcal{P}, \mathcal{T}, \mathcal{F}, I, O, M_0)$. If no two transitions have the same input/output bags, we can identify each t with the ordered pair $(I(t), O(t))$. The Petri net can then be viewed as a triple $(\mathcal{P}, \mathcal{T} \subset \mathbb{N}^{\mathcal{P}} \times \mathbb{N}^{\mathcal{P}}, M_0 \in \mathbb{N}^{\mathcal{P}})$. (In particular, the flow relation \mathcal{F} is encoded in \mathcal{T} .)

Petri nets have appeared with this presentation in different contexts and under different names: *vector addition systems* (see for instance [23]), or *chemical reaction networks* (see for instance [11, 2]).

In the chemical context, the elements of \mathcal{P} are species. The marking is the number of molecules of the different species. The elements of \mathcal{T} are *reactions*. A reaction $(c, d) \in \mathbb{N}^{\mathcal{P}} \times \mathbb{N}^{\mathcal{P}}$ is represented as follows:

$$\sum_{p \in \mathcal{P}} c_p p \longrightarrow \sum_{p \in \mathcal{P}} d_p p.$$

Example. The “chemical” form of the Petri net in Figure 1 is:



Let us now introduce two notions, deficiency and weak reversibility, which are borrowed from the chemical literature.

Definition 2.3 (Reaction graph). Let $(\mathcal{P}, \mathcal{T} \subset \mathbb{N}^{\mathcal{P}} \times \mathbb{N}^{\mathcal{P}}, M)$ be a Petri net. A complex is a vector u in $\mathbb{N}^{\mathcal{P}}$ such that: $\exists v \in \mathbb{N}^{\mathcal{P}}, (u, v) \in \mathcal{T}$ or $(v, u) \in \mathcal{T}$. The set of all complexes is denoted by \mathcal{C} . The reaction graph associated to the Petri net is the directed graph with

- nodes: \mathcal{C} , arcs: $u \rightarrow v$ if $(u, v) \in \mathcal{T}$.

Let A be the *node-arc incidence matrix* of the reaction graph, that is the $(\mathcal{C} \times \mathcal{T})$ -matrix defined by

$$A_{u,t} = -\mathbf{1}_{\{I(t)=u\}} + \mathbf{1}_{\{O(t)=u\}}. \quad (3)$$

Lemma 2.4. Consider a Petri net with set of complexes \mathcal{C} . Let ℓ be the number of connected components of the reaction graph. The rank of the node-arc incidence matrix satisfies:

$$\text{rank}(A) = |\mathcal{C}| - \ell. \quad (4)$$

Proof. Assume that $\ell = 1$. Consider $x \in \mathbb{R}^{\mathcal{C}} - \{(0, \dots, 0)\}$ such that $xA = (0, \dots, 0)$. Let C be such that $x_C \neq 0$. Consider $D \in \mathcal{C}$. Since $\ell = 1$, there exists an undirected path $(C = C_0) - C_1 - \dots - (C_k = D)$ in the reaction graph. Assume wlog that $C_i \rightarrow C_{i+1}$ and let $t_i \in \mathcal{T}$ be such that $I(t_i) = C_i, O(t_i) = C_{i+1}$. By definition of A , we have $(xA)_{t_i} = x_{C_{i+1}} - x_{C_i}$. So we have $x_{C_i} = x_{C_{i+1}}$ for all i , and $x_C = x_D$. We have proved that

$$xA = (0, \dots, 0) \implies x \in \mathbb{R}(1, \dots, 1).$$

Conversely, by definition of A , we have $(1, \dots, 1)A = (0, \dots, 0)$. We conclude that $\text{rank}(A) = |\mathcal{C}| - 1$. For a general value of ℓ , we get similarly that $\text{rank}(A) = |\mathcal{C}| - \ell$. \square

A central notion in what follows is the *deficiency* of a Petri net.

Definition 2.5 (Deficiency). *With the above notations, the deficiency of the Petri net is*

$$\delta = |\mathcal{C}| - \ell - \text{rank}(N) = \text{rank}(A) - \text{rank}(N).$$

Of particular importance are the Petri nets with **deficiency 0**. This class will be central in the study of Markovian Petri nets having a product form, see Section 3. Such Petri nets are “extremal”, in a sense made precise by next proposition.

Proposition 2.6. *For all $x \in \mathbb{R}^{\mathcal{T}}$:*

$$[Ax = 0] \implies [Nx = 0]. \quad (5)$$

In particular it implies that: $\text{rank}(A) \geq \text{rank}(N)$. Equivalently, the deficiency of a Petri net is always greater or equal to 0.

The non-negativity of the deficiency appears in Feinberg [11], using a different argument.

Proof. Using the definition of the matrix A , we can rewrite the condition $Ax = 0$:

$$\forall c \in \mathcal{C}, \quad \sum_{t: O(t)=c} x_t - \sum_{t: I(t)=c} x_t = 0, \quad (6)$$

Now for each place s , we have:

$$\begin{aligned} \sum_{t \in \mathcal{T}} N_{s,t} x_t &= \sum_{t \in \mathcal{T}} [O(t)_s - I(t)_s] x_t \\ &= \sum_{c \in \mathcal{C}} \left(\sum_{t: O(t)=c} c_s x_t - \sum_{t: I(t)=c} c_s x_t \right) = \sum_{c \in \mathcal{C}} c_s \left(\sum_{t: O(t)=c} x_t - \sum_{t: I(t)=c} x_t \right). \end{aligned}$$

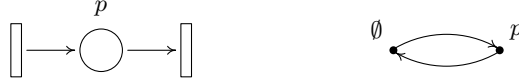
Using (6), this last sum is equal to 0.

The inequality $\text{rank}(A) \geq \text{rank}(N)$ is equivalent to $\dim \ker(A) \leq \dim \ker(N)$, which follows immediately from (5). \square

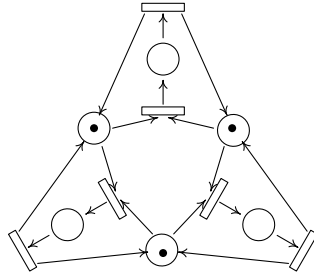
The second central notion is *weak reversibility*.

Definition 2.7 (Weak reversibility). A Petri net is weakly reversible (WR) if every connected component of the reaction graph is strongly connected.

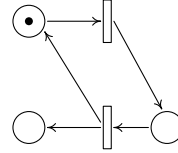
Weak reversibility is a restrictive property, see Section 4.3. It is important to observe that a connected and weakly reversible Petri net is not necessarily strongly connected. An example is given below.



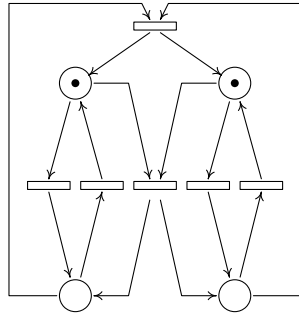
Elementary circuits of the reaction graph can be identified with the so-called “minimal closed support T-invariants” of the Petri net literature (see [5]). In particular, a Petri net is weakly reversible if and only if it is covered by minimal closed support T-invariants. Such Petri nets are called *Π -nets* in [14].



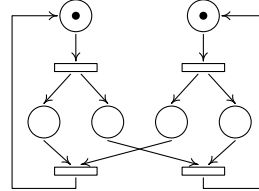
WR, $\delta = 0$



Not WR, $\delta = 0$



WR, $\delta > 0$



Not WR, $\delta > 0$

Fig. 3. Deficiency zero and weak reversibility are independent.

Weak reversibility and deficiency 0. Weak reversibility and deficiency 0 are two independent properties. The Petri net of Figure 1 is weakly reversible and has deficiency 0. If we remove one transition, t_1 for instance, weak reversibility disappears, yet the deficiency remains 0. The Petri net of Figure 4 is weakly reversible and has deficiency 1. If we remove one transition, the new Petri net is neither weakly reversible nor of deficiency 0.

Other examples of the four possible situations are given in Figure 3. The upper-left Petri net is an instance of the famous “dining philosopher” model. The Petri nets in Fig. 3 are live and bounded except for the upper-right one.

Algorithmic complexity. Weak reversibility and deficiency 0 are algorithmically simple to check. Let us determine the time complexity of the algorithms with respect to the size of the Petri net (number of places and transitions, and the unary representation of weights).

Observe first that the number of complexes is bounded by $2\mathcal{T}$. Building the reaction graph from the Petri graph can be done in time $\mathcal{O}(\mathcal{PT}^2)$. A depth-first-search algorithm on the reaction graph enables to check the weak reversibility and to compute the number of connected components (ℓ). The DFS algorithm runs in time $\mathcal{O}(\mathcal{PT})$. Computing the rank of the incidence matrix can be done in time $\mathcal{O}(\mathcal{PT}^2)$ using a Gaussian elimination.

Globally the complexity is $\mathcal{O}(\mathcal{PT}^2)$ for computing the deficiency as well as for checking weak-reversibility.

2.3 Markovian Petri nets with race policy

A Petri net is a logical object with no physical time involved. There exist several alternative ways to define timed models of Petri nets, see for instance [1, 4]. We consider the model of Markovian Petri nets with race policy. The rough description is as follows.

With each enabled transition is associated a “countdown clock” whose positive initial value is set at random. When a clock reaches 0, the corresponding transition fires. This changes the set of enabled transitions and all the clocks get reinitialized. The initial values of the clocks are chosen independently, according to an exponential distribution whose rate depends on the transition and on the current marking. With probability 1, no two clocks reach zero at the same time so the model is unambiguously defined. Enabled transitions are involved in a “race”: the transition to fire is the one whose clock will reach zero first.

We now proceed to a formal definition of the model.

When $I(t) = O(t)$, the firing of transition t does not modify the marking. For simplicity, we assume from now on that $I(t) \neq O(t)$ for all t .

Definition 2.8 (Markovian Petri net with race policy). A Markovian Petri net (with race policy) is formed by a Petri net $(\mathcal{P}, \mathcal{T}, \mathcal{F}, I, O, M_0)$ and

a set of rate functions $(\mu_t)_{t \in \mathcal{T}}$, $\mu_t : \mathcal{R}(M_0) \rightarrow \mathbb{R}_+^*$, satisfying

$$\mu_t(M) = \begin{cases} \kappa_t \Psi(M - I(t)) \Phi(M) & \text{if } M \geq I(t) \\ 0 & \text{otherwise} \end{cases}, \quad (7)$$

for some constants $\kappa_t \in \mathbb{R}_+^*$, $t \in \mathcal{T}$, and some functions Ψ and Φ valued in \mathbb{R}_+^* . The marking evolves as a continuous-time jump Markov process with state space $\mathcal{R}(M_0)$ and infinitesimal generator $Q = (q_{M,M'})_{M,M'}$, given by

$$q_{M,M'} = \sum_{t: M \xrightarrow{t} M'} \mu_t(M). \quad (8)$$

The shape (8) for the infinitesimal generator is the transcription of the informal description given at the beginning of the section.

The condition (7) for the rate functions $(\mu_t)_{t \in \mathcal{T}}$ is the same as the one in [15] and [14, Section 2]. (In [14, Section 3], an even more general shape for the rate function is considered.) Condition (7) is specifically cooked up in order for the product form result of Theorem 3.2 to hold, which explains its artificial shape. This general condition englobes two classical types of rate functions: the constant rates and the mass-action rates.

Constant rates. In the Petri net literature, the standard assumption is that the firing rates are constant:

$$\exists \kappa_t \in \mathbb{R}_+^*, \forall M \in \mathcal{R}(M_0), I(t) \geq M, \quad \mu_t(M) = \kappa_t. \quad (9)$$

Mass-action rates. In the chemical literature, the rate is often proportional to the number of different subsets of tokens (i.e. molecules) that can be involved in the firing (i.e. reaction). More precisely:

$$\forall M \in \mathcal{R}(M_0), I(t) \geq M, \quad \mu_t(M) = \kappa_t \prod_{p: I(t)_p \neq 0} \frac{M_p!}{(M_p - I(t)_p)!}. \quad (10)$$

Such rates are said to be of *mass-action* form and the corresponding stochastic process has *mass-action kinetics*. To obtain (10) from (7), set $\Phi, \Psi^{-1} : \mathbb{N}^P \rightarrow \mathbb{R}_+^*$, $x \mapsto \prod_p x_p!$.

3 Product form results

We are interested in the equilibrium behavior of Markovian Petri nets. This section presents the product form results which exist in the literature. We gather results which were spread out, obtained independently either in the Petri net community, or in the chemical one.

Let Q be the infinitesimal generator of the marking process. An invariant measure π of the process is characterized by the balance equations $\pi Q = 0$, that is: $\forall x \in \mathcal{R}(M_0)$,

$$\pi(x) \sum_{t: x \geq I(t)} \mu_t(x) = \sum_{t: x \geq O(t)} \pi(x + I(t) - O(t)) \mu_t(x + I(t) - O(t)). \quad (11)$$

A stationary distribution is an invariant probability measure. It is characterized by $\pi Q = 0$, $\sum_x \pi(x) = 1$. If π is an invariant measure and $K = \sum_{x \in \mathcal{R}(M_0)} \pi(x) < +\infty$, then $\pi/K = (\pi(x)/K)_x$ is a stationary distribution.

When the marking graph is strongly connected, the marking process is irreducible. It follows from basic Markovian theory that the stationary distribution is unique when it exists (the ergodic case). When the state space is finite, irreducibility implies ergodicity.

3.1 Non-linear traffic equations and Kelly's Theorem

Definition 3.1 (Non-linear traffic equations). Consider a Markovian Petri net with general rates. Let \mathcal{C} be the set of complexes. We call non-linear traffic equations (NLTE) the equations over the unknowns $(x_p)_{p \in \mathcal{P}}$ defined by: $\forall C \in \mathcal{C}$,

$$\prod_{p: C_p \neq 0} x_p^{C_p} \sum_{t: I(t)=C} \kappa_t = \sum_{t: O(t)=C} \kappa_t \prod_{p: I(t)_p \neq 0} x_p^{I(t)_p}. \quad (12)$$

(With the convention that the product over an empty set of indices equals 1.)

The NLTE can be viewed as a kind of balance equations (what goes in equals what goes out) at the level of complexes. Their central role appears in next theorem which is essentially due to Kelly [17, Theorem 8.1] (see also [2, Theorem 4.1]). In Kelly's book, the setting is more restrictive, but the proof carries over basically unchanged. For the sake of completeness, we recall the proof.

Theorem 3.2 (Kelly). Consider a Markovian Petri net. Assume that the NLTE (12) admit a strictly positive solution $(u_p)_{p \in \mathcal{P}}$. Then the marking process of the Petri net has an invariant measure π defined by: $\forall x \in \mathcal{R}(M_0)$,

$$\pi(x) = \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p}. \quad (13)$$

We say that π has a *product form*: $\pi(x)$ decomposes as a product over the places p of terms depending only on the local marking x_p .

Observe that $\pi(x) > 0$ for all x in (13). In particular it implies that the marking process is irreducible. On the other hand, the measure defined in (13) may have a finite or infinite mass. When it has a finite mass, the marking process is ergodic, and the normalization of π is the unique stationary distribution.

In the case of mass-action rates (10), we get

$$\sum_{x \in \mathcal{R}(M_0)} \pi(x) = \sum_{x \in \mathcal{R}(M_0)} \prod_{p \in \mathcal{P}} \frac{u_p^{x_p}}{x_p!} \leq \sum_{x \in \mathbb{N}^{\mathcal{P}}} \prod_{p \in \mathcal{P}} \frac{u_p^{x_p}}{x_p!} = \exp\left(\sum_p u_p\right) < +\infty.$$

So we are always in the ergodic case. For constant rates (9), if the state space $\mathcal{R}(M_0)$ is infinite, the ergodicity depends on the values of the constants κ_t .

Proof. It suffices to verify that π of the form (13) satisfies (11) when $(u_p)_{p \in \mathcal{P}}$ is a solution to (12).

In (11), by replacing π and μ_t with the right-hand sides of (13) and (7), we obtain, after simplification:

$$\begin{aligned} \sum_{t: x \geq I(t)} \kappa_t \Psi(x - I(t)) \Phi(x) \Phi(x)^{-1} \prod_p u_p^{x_p} &= \\ \sum_{t: x \geq O(t)} \kappa_t \Psi(x - O(t)) \Phi(x + I(t) - O(t)) \Phi(x + I(t) - O(t))^{-1} \prod_p u_p^{x_p + I(t)_p - O(t)_p} \end{aligned}$$

which is equivalent to

$$\sum_{x \geq C} \sum_{t: I(t)=C} \kappa_t \Psi(x - C) \prod_p u_p^{x_p} = \sum_{x \geq C} \sum_{t: O(t)=C} \kappa_t \Psi(x - C) \prod_p u_p^{x_p + I(t)_p - C_p}.$$

A sufficient condition for the above equality to hold is to have, for each $C \in \mathcal{C}$,

$$\sum_{t: I(t)=C} \kappa_t \Psi(x - C) \prod_p u_p^{x_p} = \sum_{t: O(t)=C} \kappa_t \Psi(x - C) \prod_p u_p^{x_p + I(t)_p - C_p},$$

after simplification we get,

$$\prod_{p: C_p \neq 0} u_p^{C_p} \sum_{t: I(t)=C} \kappa_t = \sum_{t: O(t)=C} \kappa_t \prod_{p: I(t)_p \neq 0} u_p^{I(t)_p}.$$

This last set of equations means precisely that $(u_p)_{p \in \mathcal{P}}$ is a solution to the NLTE. \square

Theorem 3.2 is the core result. Below, all the developments consist in determining conditions under which Theorem 3.2 applies. More precisely, we want conditions on the model ensuring the existence of a strictly positive solution to the NLTE and the finiteness of the measure π . The ideal situation is as follows:

- *structural* properties of the Petri net (i.e. independent of the firing rates) ensure the existence of a strictly positive solution to the NLTE;
- conditions on the firing rates ensure the finiteness of the measure π .

3.2 Linear traffic equations and Haddad & al's Theorem

Solving the non-linear traffic equations is still a challenging task. We may avoid a direct attack to these equations by considering a simpler system of equations called the linear traffic equations.

Definition 3.3 (Linear traffic equations). *We call linear traffic equations (LTE) the equations over the unknowns $(y_C)_{C \in \mathcal{C}}$ defined by: $\forall C \in \mathcal{C}$,*

$$y_C \sum_{t: I(t)=C} \kappa_t = \sum_{t: O(t)=C} \kappa_t y_{I(t)}. \quad (14)$$

Furthermore, if $\emptyset \in \mathcal{C}$, then $y_\emptyset = 1$.

The NLTE and the LTE are clearly linked.

Lemma 3.4. *If the NLTE (12) have a strictly positive solution $u = (u_p)_{p \in \mathcal{P}}$, then $v = (v_C)_{C \in \mathcal{C}}$,*

$$v_C = \prod_{p: C_p \neq \emptyset} u_p^{C_p}, \quad (15)$$

is a strictly positive solution to the LTE (14).

For a partial converse statement, see Lemma 3.10. The following proposition provides a simple and structural criterium for the existence of a strictly positive solution to the LTE.

Proposition 3.5. *The following statements are equivalent:*

- $\exists (\kappa_t)_{t \in \mathcal{T}}$ such that the equations (14) have a strictly positive solution.
- $\forall (\kappa_t)_{t \in \mathcal{T}}$, the equations (14) have a strictly positive solution.
- The Petri net is weakly reversible.

Proofs can be found in [5, Theorem 3.5] or [11, Corollary 4.2]. We recall the argument from [5] which is simple and illuminating.

Proof. The *reaction process* is a continuous-time Markov process, analog to the marking process, except that it is built on the reaction graph instead of the marking graph. More precisely, the state space is the set of complexes \mathcal{C} and the infinitesimal generator $\tilde{Q} = (\tilde{q}_{u,v})_{u,v}$ is defined by

$$\tilde{q}_{u,v} = \sum_{t: I(t)=u, O(t)=v} \kappa_t.$$

(The discrete-time version of this process was introduced in [15] under the name “routing process”.) The key observation is that the LTE (14) are precisely the balance equations $y\tilde{Q} = 0$ of the reaction process. The result now follows using standard Perron-Frobenius theory. \square

So weak reversibility is a necessary condition to have a strictly positive solution to the NLTE, and to be able to apply Theorem 3.2. Unfortunately, it is not a sufficient condition as shown by the following example.

Example. Let us consider a Markovian Petri net whose underlying Petri graph is shown in Figure 4, and is equivalently defined by the chemical reactions:



This is a weakly reversible Petri net, thus its LTE always have a strictly positive solution regardless of the choice of the constants κ_t . The NLTE are:

$$\kappa_1 x_1 = \kappa_2 x_2 \quad \kappa_3 x_3 = \kappa_4 x_4 \quad \kappa_5 x_1 x_3 = \kappa_6 x_2 x_4 . \quad (16)$$

The system (16) does not always have a strictly positive solution. For example, set $\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = \kappa_5 = 1$, and $\kappa_6 = 2$. Any solution to (16) must satisfy either $x_1 = x_2 = 0$ or $x_3 = x_4 = 0$.

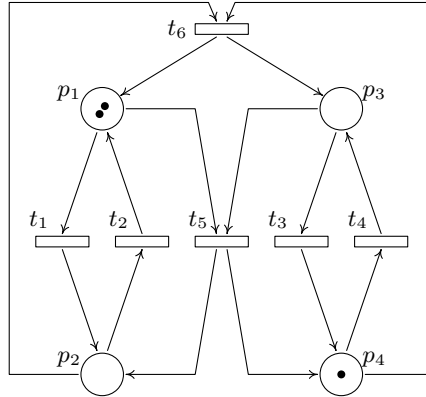


Fig. 4. A weakly reversible Petri net.

Depending on the values of the constants $(\kappa_t)_t$, the Markovian Petri net may or may not have a product form invariant measure. Anticipating on Theorem 3.8, the deficiency of the Petri net has to be different from 0, and indeed we have a deficiency which is equal to 1.

So now the goal is to find additional conditions on top of weak reversibility to ensure the existence of a product form.

An early result in this direction appears in Coleman, Henderson and Taylor [6, Theorem 3.1]. The condition is not structural (i.e. rate dependent) and not very tractable. Next result, due to Haddad, Moreaux, Sereno, and Silva [14, Theorem 9], provides a structural sufficient condition.

Proposition 3.6. *Consider a Markovian Petri net (set of complexes \mathcal{C}). Assume that the Petri net is weakly reversible. Let N be the incidence matrix of the Petri net, see (1). Let A be the node-arc incidence matrix of the reaction graph, see (3). Assume that there exists a \mathbb{Q} -valued $(\mathcal{C} \times \mathcal{P})$ -matrix B such that $BN = A$. Then the marking process has an invariant measure π given by: $\forall x \in \mathcal{R}(M_0)$,*

$$\pi(x) = \Phi(x)^{-1} \prod_{p \in \mathcal{P}} \left(\prod_{C \in \mathcal{C}} v_C^{B_{C,p}} \right)^{x_p},$$

where v is a strictly positive solution to the LTE.

3.3 Deficiency zero and product form

Independently of the efforts in the Petri net community ([6, 14]), the following result was proved on the chemical side by Feinberg [11, Theorem 5.1].

Theorem 3.7 (Feinberg). *Consider a Markovian Petri net. Assume that the Petri net has deficiency 0. Then the NLTE have a strictly positive solution if and only if the network is weakly reversible.*

By combining Theorems 3.2 and 3.7, we obtain the following result whose formulation is original.

Theorem 3.8. *Consider a Petri net which is weakly reversible and has deficiency 0. Consider any associated Markovian Petri net. The NLTE have a strictly positive solution $(u_p)_p$ and the marking process has a product form invariant measure:*

$$\pi(x) = \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p}. \quad (17)$$

If we assume furthermore that the rates are of mass-action type (10), then the marking process is ergodic and its stationary distribution is:

$$\pi(x) = C \prod_{p \in \mathcal{P}} \frac{u_p^{x_p}}{x_p!},$$

where $C = (\sum_x u_p^{x_p} / x_p!)^{-1}$.

The above result is interesting. Indeed, the “deficiency 0” condition is structural and very simple to handle. We now prove that the result in Theorem 3.8 is equivalent to the one in Proposition 3.6.

Proposition 3.9. *Consider a Petri net. There exists a $(\mathcal{C} \times \mathcal{P})$ -matrix B such that $BN = A$ (with the notations of Prop. 3.6) if and only if the Petri net has deficiency 0.*

Proof. The deficiency of the Petri net is 0 iff $\text{rank}(N) = \text{rank}(A)$.

Assume first that there exists a matrix B such that $BN = A$. Since $BN = A$, we have $\text{rank}(A) \leq \text{rank}(N)$. By Proposition 2.6, we also have $\text{rank}(A) \geq \text{rank}(N)$. Therefore the Petri net has deficiency 0.

We now prove the converse result. Assume that the Petri net has deficiency 0. Set $r = |\mathcal{C}| - \ell$. We have $\text{rank}(A) = \text{rank}(N) = r$. Since $\text{rank}(A) = r$, we know that there exists an invertible and \mathbb{Q} -valued $(\mathcal{T} \times \mathcal{T})$ -matrix Q such that the first r column vectors of AQ are linearly independent and the last $(|\mathcal{T}| - r)$ column vectors are $(0, \dots, 0)^T$. According to (5), the last $(|\mathcal{T}| - r)$ column vectors of NQ are $(0, \dots, 0)^T$. But $\text{rank}(N) = r$, so the first r column vectors of NQ must be linearly independent.

Denote by AQ_1, \dots, AQ_r , *resp.* NQ_1, \dots, NQ_r , the first r column vectors of AQ , *resp.* NQ . Since the two families are independent, we know that there exists a \mathbb{Q} -valued $(\mathcal{C} \times \mathcal{P})$ -matrix B such that $BNQ_i = AQ_i$, for all $i = 1, \dots, r$. In other words,

$$BNQ = AQ. \quad (18)$$

Finally, right-multiplying both sides of (18) by Q^{-1} , we obtain $BN = A$. \square

The main results are summarized in the diagram of Figure 5.

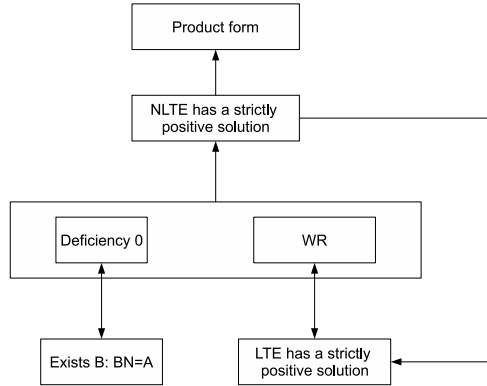


Fig. 5. Deficiency 0, weak reversibility, and product form

3.4 A detailed example

Consider the Petri graph represented on the left of Figure 6. The corresponding reaction graph is given on the right of the figure.

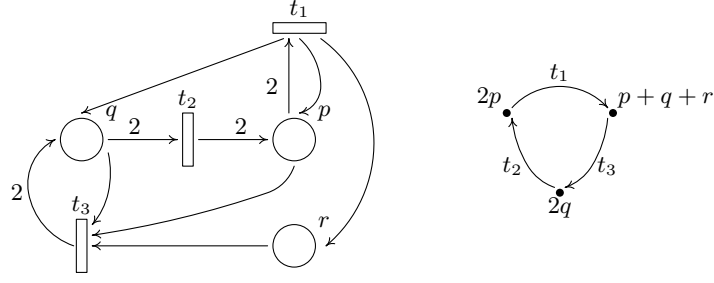


Fig. 6. A Petri graph and its reaction graph.

The reaction graph is strongly connected so the Petri graph is weakly reversible. The incidence matrices N and A are given by (indices are ranked as (p, q, r) , (t_1, t_2, t_3) , and $(2p, p + q + r, 2q)$):

$$N = \begin{pmatrix} -1 & 2 & -1 \\ 1 & -2 & 1 \\ 1 & 0 & -1 \end{pmatrix}, \quad A = \begin{pmatrix} -1 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}.$$

We check that $\text{rank}(A) = \text{rank}(N) = 2$, so the deficiency is 0. We are in the scope of application of the results of Section 3.3.

Denote a marking M by the triple (M_p, M_q, M_r) . Consider the two Petri nets corresponding to the above Petri graph with two different initial markings: $(2, 0, 0)$ and $(3, 0, 0)$.

The two Petri nets have a drastically different behaviour. The first one is live and bounded, while the second one is live and unbounded. The sets of reachable markings are, respectively,

$$\begin{aligned} \mathcal{R}(2, 0, 0) &= \{(2, 0, 0), (1, 1, 1), (0, 2, 0)\} \\ \mathcal{R}(3, 0, 0) &= \{(1, 2, 0), (0, 3, 1)\} \cup \{(i, 3 - i, 2n + 1 - i), 0 \leq i \leq 3, n \geq 1\} \end{aligned}$$

For the initial marking $(2, 0, 0)$, the marking graph is the elementary circuit $(2, 0, 0) \rightarrow (1, 1, 1) \rightarrow (0, 2, 0) \rightarrow (2, 0, 0)$. For the initial marking $(3, 0, 0)$, the marking graph is represented in Figure 7. The dashed arrows correspond to transition t_1 , the dash-and-dotted ones to t_2 , and the plain ones to t_3 .

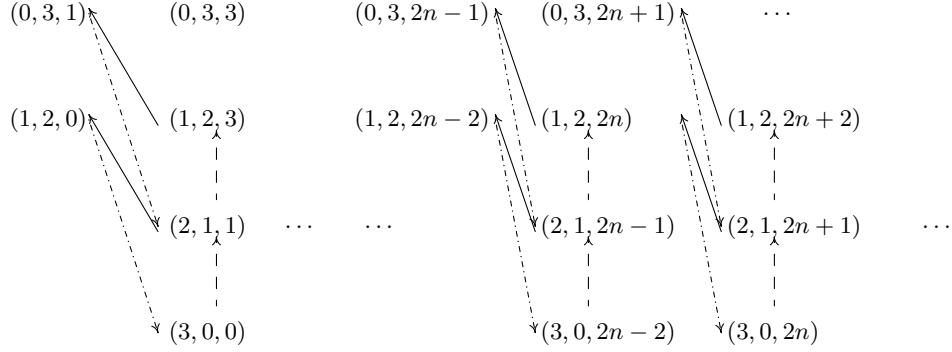


Fig. 7. Marking graph with the initial marking $(3, 0, 0)$.

Consider the Markovian Petri nets associated with the above Petri nets and constant firing rates $(\kappa_1, \kappa_2, \kappa_3)$ for (t_1, t_2, t_3) . The NLTE over the unknowns (x_p, x_q, x_r) are given by

$$\kappa_1 x_p^2 = \kappa_2 x_q^2, \quad \kappa_3 x_p x_q x_r = \kappa_1 x_p^2, \quad \kappa_2 x_q^2 = \kappa_3 x_p x_q x_r.$$

A strictly positive solution to the NLTE is

$$\left(\frac{\sqrt{\kappa_2}}{\sqrt{\kappa_1}}, 1, \frac{\sqrt{\kappa_1 \kappa_2}}{\kappa_3} \right).$$

Let \mathcal{R} denote the set of reachable markings. According to Theorem 3.8, the invariant measure π is given by

$$\forall m = (m_p, m_q, m_r) \in \mathcal{R}, \quad \pi(m) = \kappa_1^{(m_r - m_p)/2} \kappa_2^{(m_p + m_r)/2} \kappa_3^{-m_r}. \quad (19)$$

The invariant measure is expressed in exactly the same way for the two Petri nets. But it corresponds to two very different situations.

For the initial marking $(2, 0, 0)$, we have $|\mathcal{R}| = 3$, the model is ergodic and the unique stationary distribution p , obtained by normalization of (19), is given by

$$p(2, 0, 0) = C^{-1} \kappa_2 \kappa_3, \quad p(1, 1, 1) = C^{-1} \kappa_1 \kappa_2, \quad p(0, 2, 0) = C^{-1} \kappa_1 \kappa_3,$$

with $C = \kappa_1 \kappa_2 + \kappa_2 \kappa_3 + \kappa_1 \kappa_3$.

For the initial marking $(3, 0, 0)$, we have $|\mathcal{R}| = \infty$, and the model is ergodic if and only if the following stability condition is satisfied

$$\kappa_1 \kappa_2 < \kappa_3^2.$$

It is interesting to observe that we get a non-linear stability condition.

3.5 Additional results

The statements of Proposition 3.6 and Theorem 3.8 differ in that they rely respectively on the LTE and the NLTE. In particular, it seems at first glance that Prop. 3.6 manages to bypass the NLTE. But it is not the case: the NLTE are hidden in the matrix B , see below.

Lemma 3.10. *Consider a weakly reversible Markovian Petri net. Assume that there exists a \mathbb{Q} -valued matrix B such that $BN = A$ (notations of Prop. 3.6). Let $v = (v_C)_{C \in \mathcal{C}}$ be a strictly positive solution to the LTE. Then $u = (u_p)_{p \in \mathcal{P}}, u_p = \prod_{C \in \mathcal{C}} v_C^{B_{C,p}}$ is a strictly positive solution to the NLTE.*

Proof. For each transition t , we have:

$$\begin{aligned} \prod_p u_p^{I(t)_p - O(t)_p} &= \prod_p u_p^{-N_{p,t}} = \prod_p \left(\prod_C v_C^{B_{C,p}} \right)^{-N_{p,t}} \\ &= \prod_C v_C^{-\sum_p B_{C,p} N_{p,t}} = \prod_C v_C^{-A_{C,t}} \\ &= \frac{v_{I(t)}}{v_{O(t)}}. \end{aligned}$$

Then we have

$$\begin{aligned} [u \text{ strictly positive sol. NLTE}] &\iff \forall C, \sum_{I(t)=C} \kappa_t = \sum_{O(t)=C} \kappa_t \prod_p u_p^{I(t)_p - O(t)_p} \\ &\iff \forall C, \sum_{I(t)=C} \kappa_t = \sum_{O(t)=C} \kappa_t \frac{v_{I(t)}}{v_{O(t)}} \\ &\iff [v \text{ strictly positive sol. LTE}]. \end{aligned}$$

The proof is completed. \square

Let us comment on a specific point. Consider Theorem 3.8. The invariant measure is defined in fonction of a specific strictly positive solution to the NLTE. However, it is easily seen that the NLTE may have several strictly positive solutions. Is this contradictory with the uniqueness of the stationary measure in the ergodic case ? In the non-ergodic case, do we get several invariant measures ? Next result answers these questions.

Lemma 3.11. *Assume that the Petri net is weakly reversible and has deficiency 0. Let u, \tilde{u} be solutions to the NLTE and $\pi, \tilde{\pi}$ be the corresponding invariant measures (given by (17)). Then there exists a constant K such that for all reachable marking x , $\tilde{\pi}(x) = K\pi(x)$.*

Proof. It suffices to show that

$$\frac{\pi(x - I(t) + O(t))}{\pi(x)} = \frac{\tilde{\pi}(x - I(t) + O(t))}{\tilde{\pi}(x)}, \quad (20)$$

for all reachable marking x and for all enabled transition t of x . Define $v = (v_C)_{C \in \mathcal{C}}$, and $\tilde{v} = (\tilde{v}_C)_{C \in \mathcal{C}}$ by

$$v_C = \prod_p u_p^{C_p}, \quad \tilde{v}_C = \prod_p \tilde{u}_p^{C_p}.$$

According to Lemma 3.4, v and \tilde{v} are solutions to the LTE. Equality (20) holds if and only if

$$\prod_p u_p^{O(t)_p - I(t)_p} = \prod_p \tilde{u}_p^{O(t)_p - I(t)_p} \iff \frac{v_{O(t)}}{v_{I(t)}} = \frac{\tilde{v}_{O(t)}}{\tilde{v}_{I(t)}} \iff \frac{\tilde{v}_{O(t)}}{v_{O(t)}} = \frac{\tilde{v}_{I(t)}}{v_{I(t)}}.$$

The last equality is proved by Feinberg in [11, Proposition 4.1]. \square

3.6 Algorithmic complexity

Let us compare Proposition 3.6 and Theorem 3.8 from an algorithmic point of view.

In both cases, one needs to check weak reversibility. Using Proposition 3.5, weak reversibility is equivalent to the existence of a strictly positive solution to the LTE. This last point can be checked in time $\mathcal{O}(\mathcal{C}^3)$. Then the procedures diverge.

- *Proposition 3.6.* We need to compute the matrix B satisfying $BN = A$. This requires to solve \mathcal{C} linear systems of dimension $\mathcal{P} \times \mathcal{T}$. So the time-complexity is $\mathcal{O}(\mathcal{CPT}^2)$ using Gaussian elimination.
- *Theorem 3.8.* We have seen in Section 2.2 that the deficiency 0 condition can be checked in time $\mathcal{O}(\mathcal{PT}^2)$. Then one needs to compute a strictly positive solution to the NLTE. This can be done as follows. Consider (15), apply the logarithm operation on both sides and solve the linear system. The corresponding time-complexity is $\mathcal{O}(\mathcal{PC}^2)$ ($= \mathcal{O}(\mathcal{PT}^2)$ since $|\mathcal{C}| \leq 2|\mathcal{T}|$).

We conclude that it is more efficient to determine the product form by using the characterization in Theorem 3.8.

4 Markovian free-choice nets and product form

The class of Petri nets whose Markovian version have a product form is an interesting one. It is therefore natural to study how this class intersects with the classical families of Petri nets: state machines and free-choice Petri nets.

The central result of this section is, in a sense, a negative result. We show that within the class of free-choice Petri nets, the only ones which are weakly reversible are closely related to state machines. We also show that the Markovian state machines are “equivalent to” Jackson networks. The latter form the most basic and classical example of product form queueing networks.

From now on, we consider only *non-weighted* Petri nets, that is Petri nets with $I, O : \mathcal{T} \rightarrow \{0, 1\}^{\mathcal{P}}$. In this case, the input/output bags can be retrieved from the flow relation \mathcal{F} and we can define the Petri net as a quadruple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$. We also identify complexes with subsets of \mathcal{P} .

For a node $x \in \mathcal{T} \cup \mathcal{P}$, set $\bullet x = \{y : (y, x) \in \mathcal{F}\}$ and $x^\bullet = \{y : (x, y) \in \mathcal{F}\}$. For a set of nodes $S \subset \mathcal{T} \cup \mathcal{P}$, set $\bullet S = \bigcup_{x \in S} \bullet x$ and $S^\bullet = \bigcup_{x \in S} x^\bullet$.

4.1 State machines

Definition 4.1 (State machine and generalized state machine). A *non-weighted Petri net* $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$ is a:

- State machine (SM) if for all transition t , $|\bullet t| = |t^\bullet| = 1$;
- Generalized state machine (GSM) if for all transition t , $|\bullet t| \leq 1$, $|t^\bullet| \leq 1$.

Definition 4.2 (Associated state machine). Given a generalized state machine $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$, the associated state machine is $\mathcal{N}' = (\mathcal{P}', \mathcal{T}, \mathcal{F}', M'_0)$ in which:

- $\mathcal{P}' = \mathcal{P} \cup \{p\}, p \notin \mathcal{P}$,
- $\mathcal{F}' = \mathcal{F} \cup \{(p, t), t \in \mathcal{T}, |\bullet t| = 0\} \cup \{(t, p), t \in \mathcal{T}, |t^\bullet| = 0\}$,
- $\forall x \in \mathcal{P}, M'_0(x) = M_0(x), \quad M'_0(p) = 0$.

Figure 8 shows a SM, a GSM and its associated SM.

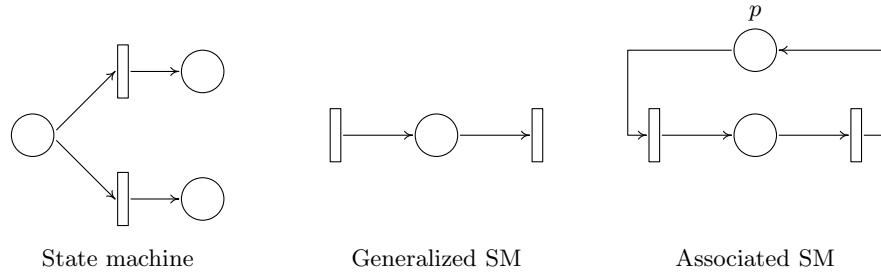


Fig. 8. State machine, generalized state machine and associated state machine.

Lemma 4.3. The reaction graph and the Petri graph of a state machine are isomorphic. The reaction graph of a GSM and the Petri graph of its associated SM are isomorphic.

Proof. In a SM, each complex is just one place. Starting from the Petri graph and replacing $[p \rightarrow t \rightarrow q]$, $p, q \in \mathcal{P}$, $t \in \mathcal{T}$, by $[p \rightarrow q]$, we get the reaction graph. For GSM, the mapping is the same with the empty complex corresponding to the “new” place in the associated SM. \square

Corollary 4.4. *A SM is weakly reversible iff each connected component is strongly connected. A GSM is weakly reversible iff in the associated SM, each connected component is strongly connected.*

In a SM, the complexes are the places. So the NLTE and the LTE coincide exactly. For a GSM, the complexes are the places and the empty set. With the convention $y_\emptyset = 1$, we still have that the NLTE and the LTE coincide. Next proposition follows.

Proposition 4.5. *Consider a weakly reversible GSM. For every rates $(\kappa_t)_t$, the NLTE have a strictly positive solution.*

Proof. In the weakly reversible case, the LTE have a strictly positive solution for every choice of the rates, Proposition 3.5. Therefore the NLTE have a strictly positive solution for every choice of the rates. \square

The above proof does not require Feinberg's Theorem 3.7. However, it turns out that the deficiency is 0, which provides a second proof of Prop. 4.5 using Theorem 3.7.

Proposition 4.6. *Generalized state machines have deficiency 0.*

Proof. Consider first a state machine. Let N be the incidence matrix, and let A be the node-arc incidence matrix of the reaction graph, see (3). Using Lemma 4.3, we get immediately that $A = N$. So, in particular, we have $\text{rank}(A) = \text{rank}(N)$ and the deficiency is 0.

Consider now a GSM \mathcal{N} and its associated SM \mathcal{N}' . Call \mathcal{C} (resp. \mathcal{C}'), N (resp. N') and ℓ (resp. ℓ') the set of complexes, the incidence matrix and the number of connected components of the reaction graph of \mathcal{N} (resp. \mathcal{N}').

Since \mathcal{N} and \mathcal{N}' have the same reaction graph (Lemma 4.3), we have:

$$|\mathcal{C}| = |\mathcal{C}'|, \ell = \ell'. \quad (21)$$

By construction of \mathcal{N}' , N' is N augmented with a row $(x_t)_{t \in \mathcal{T}}$ defined by

$$x_t = \mathbf{1}_{\{t^\bullet = \emptyset\}} - \mathbf{1}_{\{\bullet t = \emptyset\}},$$

(where t^\bullet and $\bullet t$ are defined in \mathcal{N}). We have $\text{rank}(N') \geq \text{rank}(N)$. On the other hand, observe that $\forall t \in \mathcal{T}, x_t = -\sum_{s \in \mathcal{P}} N_{s,t}$, so $N' = BN$, where B is the $\mathcal{P} \times \mathcal{P}$ identity matrix augmented with the row $(-1, \dots, -1)$. Hence $\text{rank}(N') = \text{rank}(BN) \leq \text{rank}(N)$. So:

$$\text{rank}(N') = \text{rank}(N). \quad (22)$$

Together (21) and (22) imply that \mathcal{N} and \mathcal{N}' have the same deficiency. Since \mathcal{N}' is a SM, it has deficiency zero, so \mathcal{N} also has deficiency zero. \square

By coupling Proposition 4.5 and Theorem 3.2, or alternatively Proposition 4.6 and Theorem 3.8, we get the result below.

Corollary 4.7. *Consider a Markovian weakly reversible GSM. The NLTE have a strictly positive solution $(u_p)_p$. The marking process admits a product form invariant measure given by: $\forall x \in \mathcal{R}(M_0)$,*

$$\pi(x) = \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p}.$$

In the case of a SM, $\mathcal{R}(M_0)$ is finite, the marking process is ergodic, and π can be normalized to give a product form stationary distribution: $\forall x \in \mathcal{R}(M_0)$,

$$\tilde{\pi}(x) = B\Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p},$$

where $B = (\sum_{x \in \mathcal{R}(M_0)} \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p})^{-1}$.

Corollary 4.7 is far from a surprising or new result, as we now show.

4.2 Jackson networks

The product form result for Jackson networks is one of the cornerstones of Markovian queueing theory. It was originally proved by Jackson [16] for open networks and by Gordon & Newell [13] for closed networks.

Consider a Markovian weakly reversible SM with constant rates $(\kappa_t)_{t \in \mathcal{T}}$. It can be transformed into a Jackson network as follows:

- A place s becomes a simple queue, that is a single server Markovian queue with an infinite buffer. The service rate at queue s is $\mu_s = \sum_{t \in s^\bullet} \kappa_t$.
- The routing matrix P of the Jackson network is the stochastic matrix defined as follows: $\forall u, v \in \mathcal{P}$,

$$P_{u,v} = \begin{cases} \mu_u^{-1} \sum_{t: \bullet t = u, t^\bullet = v} \kappa_t & \text{if } \exists t \in \mathcal{T}, u \rightarrow t \rightarrow v \\ 0 & \text{otherwise} \end{cases}.$$

- A token in place s becomes a customer in queue s .

Consider now a Markovian weakly reversible GSM with constant rates $(\kappa_t)_{t \in \mathcal{T}}$. On top of the above transformations, we do the following:

- A transition t with $\bullet t = \emptyset$ becomes an external Poisson arrival flow of rate κ_t in queue t^\bullet .

The routing matrix P is now substochastic. Indeed, if the transition t is such that $t^\bullet = \emptyset$, then $\sum_v P_{t^\bullet, v} < 1$.

In the SM case, the Jackson network is *closed*, that is without arrivals from the outside and without departures to the outside. In the GSM case with input and output transitions, the Jackson network is *open*.

The transformation from (G)SM to Jackson network is illustrated on Fig. 9.

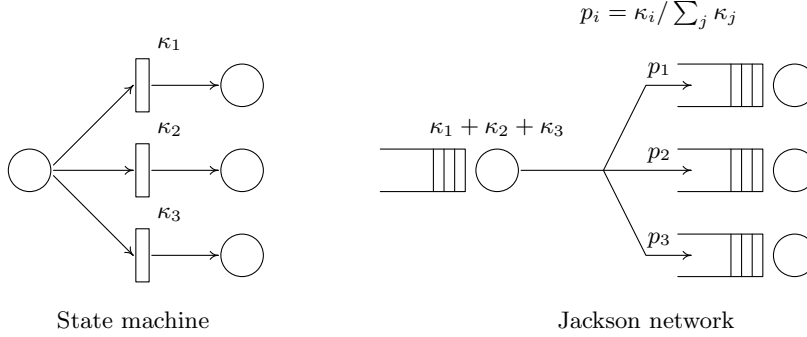


Fig. 9. From (generalized) state machine to Jackson network.

A Jackson network can be translated into a Markovian (G)SM using the same construction in the reverse direction.

The two models are identical in a strong sense. Precisely, the marking process of the state machine and the queue-length process of the Jackson network have the same infinitesimal generator.

The classical product form results for Jackson networks (Jackson [16] and Gordon & Newell [13]) are exactly the translation via the above transformation of Corollary 4.7. In the open case, the weak-reversibility implies the classical “without capture” condition of Jackson networks.

The above transformation from GSM to queueing network can also be performed in the case of general rate functions of type (7). Queueing networks with those rate functions are called *Whittle networks* in the literature. The existence of product form invariant measures for these networks is a classical result, see for instance [24] and the references therein.

4.3 Free-choice Petri nets

We study the family of live and bounded free-choice nets. This is an important class of Petri nets realizing a nice compromise between modelling power and tractability, see the dedicated monography of Desel & Esparza [8]. We show that the only such Petri nets having a product form are, in a sense, the GSM.

Definition 4.8 (Free-choice Petri net). *A free-choice Petri net is a non-weighted Petri net $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$ such that: for every two transitions t_1 and t_2 , either $\bullet t_1 = \bullet t_2$ or $\bullet t_1 \cap \bullet t_2 = \emptyset$.*

Some authors call the above an *extended free-choice* Petri net and have a more restrictive definition for free-choice Petri nets.

In Figure 10, the Petri net on the left is free-choice, while the one on the right is not free-choice.

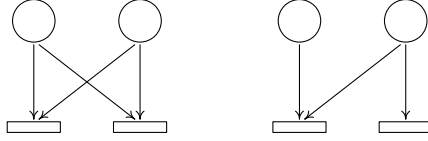


Fig. 10. Free-choice (left) and non free-choice (right) Petri nets.

Definition 4.9 (Cluster). *The cluster of a node $x \in \mathcal{P} \cup \mathcal{T}$, denoted by $[x]$, is the minimal set of nodes such that: (i) $x \in [x]$; (ii) $\forall t \in \mathcal{T} : t \in [x] \implies {}^\bullet t \subset [x]$; (iii) $\forall p \in \mathcal{P} : p \in [x] \implies p^\bullet \subset [x]$.*

The clusters form a partition of the set of nodes, see [8, Proposition 4.5], and therefore of the places. Moreover, we have the following.

Lemma 4.10. *Consider a weakly reversible free-choice Petri net. The non-empty complexes are disjoint subsets of \mathcal{P} . The partition of \mathcal{P} induced by the non-empty complexes is the same as the partition of \mathcal{P} induced by the clusters.*

Proof. In a weakly reversible free-choice Petri net, the non-empty complexes are also non-empty input bags, which are disjoint according to the definition of free-choiceness.

It follows from the definition of clusters that every non-empty input bag is entirely contained in a cluster. This cluster is unique because the clusters partition the set of places. Let I be a non-empty complex (which is also a non-empty input bag). Denote by $[I]$ the cluster containing I . We have $I^\bullet \subset [I]$, so $I \cup I^\bullet \subset [I]$. Since the Petri net is free-choice, ${}^\bullet t = I$ for all $t \in I^\bullet$. The set $I \cup I^\bullet$ satisfies the three conditions of the definition of clusters, so we have $[I] \subset I \cup I^\bullet$. We conclude that $[I] = I \cup I^\bullet$ and I is the set of places of the cluster $[I]$.

Conversely, let $[x]$ be a cluster such that $[x] \cap \mathcal{P} \neq \emptyset$. Let I be a non-empty input bag contained in $[x]$. We have, using the above, $[I] = I \cup I^\bullet \subset [x]$. By minimality, $[I] = [x]$ and $[x] \cap \mathcal{P} = I$. \square

The above result is not true for a non-WR free-choice Petri net. Consider for instance the Petri net represented in the upper-right part of Figure 3.

Under the assumptions of Lemma 4.10, the non-empty complexes are disjoint. Thus each non-empty complex behaves as if it was a “big place”. Consider the operation which reduces each non-empty complex to a single place. The resulting Petri net is a generalized state machine. And this generalized state machine is weakly reversible because the original free-choice Petri net was weakly reversible. Let us define all this more formally.

Definition 4.11 (Reduced generalized state machine). Let $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$ be a weakly reversible free-choice Petri net with set of complexes \mathcal{C} . We call the reduced generalized state machine (RGSM) of \mathcal{N} the GSM $\mathcal{RN} = (\mathcal{C} \setminus \{\emptyset\}, \mathcal{T}, \tilde{\mathcal{F}}, \tilde{M}_0)$ where:

- $\tilde{\mathcal{F}} = \{(\bullet t, t), (t, t\bullet), t \in \mathcal{T}\};$
- \tilde{M}_0 is defined by: $\forall C \in \mathcal{C} \setminus \{\emptyset\}, \tilde{M}_0(C) = \min_{p \in C} M_0(p).$

The Petri graph of \mathcal{RN} is roughly the reaction graph of \mathcal{N} reinterpreted as a Petri graph.

Lemma 4.12. Let \mathcal{N} be a weakly reversible free-choice Petri net and \mathcal{RN} its RGSM. The marking graph of \mathcal{RN} is isomorphic to the one of \mathcal{N} . If \mathcal{N} and \mathcal{RN} are Markovian with the same rates then the two marking processes are “identical”, meaning that they have the same infinitesimal generator.

Proof. Consider two places p and q belonging to the same complex. Since the complexes are disjoint, Lemma 4.10, each time p gains (resp. loses) a token, so does q . So the difference $M_p - M_q$ is invariant over all the reachable markings M . It has the following consequence.

Consider $f : \mathcal{R}(M_0) \rightarrow \mathbb{N}^{\mathcal{C} \setminus \emptyset}$ defined by $f(M)_C = \min_{p \in C} M(p)$. If $M \xrightarrow{t} M'$ in \mathcal{N} then $f(M) \xrightarrow{t} f(M')$ in \mathcal{RN} . So $f(\mathcal{R}(M_0)) = \mathcal{R}(\tilde{M}_0)$ and the marking graph of \mathcal{RN} is the marking graph of \mathcal{N} up to a renaming of the nodes.

Since the marking graphs are the same, the infinitesimal generators are also identical if the two Petri nets have the same rates. \square

One could introduce the RGSM associated with a non-WR free-choice Petri net as in Definition 4.11. But in this case Lemma 4.12 does not hold, and the two marking graphs have nothing in common. This is illustrated in Figure 11.

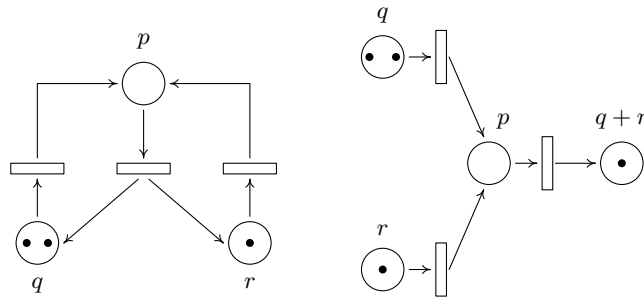


Fig. 11. A non-WR free-choice net and the associated RGSM.

Now let us compare the structural characteristics of the original free-choice Petri net \mathcal{N} and of the reduced generalized state machine \mathcal{RN} .

Lemma 4.13. *Let \mathcal{N} be a weakly reversible free-choice Petri net. The RGSM \mathcal{RN} is weakly reversible and has the same deficiency as \mathcal{N} .*

Proof. The weak reversibility of \mathcal{RN} follows directly from the definition of reduced generalized state machines.

The Petri graph of \mathcal{RN} is isomorphic to its reaction graph, Lemma 4.3. Now by construction, \mathcal{N} and \mathcal{RN} have the same reaction graph. So the number of complexes and the number of connected components of the reaction graph do not change. Call N and N' the incidence matrices of \mathcal{N} and \mathcal{RN} respectively. Let C be an arbitrary complex, let p, p' be two places of C . For every transition t , we have $N_{p,t} = N_{p',t} = N'_{C,t}$, which implies that $\text{rank}(N) = \text{rank}(N')$. So the two Petri nets have the same deficiency. \square

Corollary 4.14. *Weakly reversible free-choice Petri nets have deficiency 0.*

Proof. This follows from Prop. 4.6 and Lemma 4.13. \square

Now all the results for weakly reversible GSM can be applied to weakly reversible free-choice Petri nets. We get the following.

Theorem 4.15. *Let \mathcal{N} be a Markovian free-choice Petri net. Then \mathcal{N} is weakly reversible if and only if its NLTE have a strictly positive solution.*

In this case, the Petri net has deficiency zero. Let $(u_p)_p$ be a strictly positive solution to the NLTE. The marking process has a product form invariant measure π given by: $\forall x \in \mathcal{R}(M_0)$,

$$\pi(x) = \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p}.$$

If $\emptyset \notin \mathcal{C}$ then the state space $\mathcal{R}(M_0)$ is finite, the marking process is ergodic and π can be normalized to give a product form stationary distribution: $\forall x \in \mathcal{R}(M_0)$,

$$\tilde{\pi}(x) = B\Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p},$$

where $B = (\sum_{x \in \mathcal{R}(M_0)} \Phi(x)^{-1} \prod_{p \in \mathcal{P}} u_p^{x_p})^{-1}$.

Acknowledgement. We are grateful to Serge Haddad for several fruitful discussions and for suggesting the example of Section 3.4.

References

- [1] M. Ajmone-Marsan, G. Balbo, A. Bobbio, G. Chiola, G. Conte, and A. Cumani. The effect of execution policies on the semantics and analysis of stochastic Petri nets. *IEEE Trans. on Software Engin.*, 15(7):832–846, 1989.

- [2] D. Anderson, G. Craciun, and T. Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 2009. To appear. Preprint found at oai:arXiv.org:0803.3042.
- [3] D. Angeli, P. De Leenheer, and E. D. Sontag. A Petri net approach to the study of persistence in chemical reaction networks. *Mathematical Biosciences*, 210(2):598 – 618, 2007.
- [4] F. Baccelli, G. Cohen, G.J. Olsder, and J.P. Quadrat. *Synchronization and Linearity*. John Wiley & Sons, New York, 1992.
- [5] R. Boucherie and M. Sereno. On closed support T -invariants and the traffic equations. *J. Appl. Probab.*, 35(2):473–481, 1998.
- [6] J. L. Coleman, W. Henderson, and P. G. Taylor. Product form equilibrium distributions and a convolution algorithm for stochastic Petri nets. *Performance Evaluation*, 26(3):159 – 180, 1996.
- [7] J. Desel. Basic linear algebraic techniques for place or transition nets. In *Lectures on Petri Nets I*, number 1491 in LNCS, pages 257–308. Springer-Verlag, 1996.
- [8] J. Desel and J. Esparza. *Free Choice Petri Nets*, volume 40 of *Cambridge Tracts Theoret. Comput. Sci.* Cambridge Univ. Press, 1995.
- [9] N. M. Van Dijk. *Queueing Networks and Product Forms: A Systems Approach*. John Wiley & Sons, 1993.
- [10] S. Donatelli and M. Sereno. On the product form solution for stochastic Petri nets. In *13-th Int. Conf. on Application and Theory of Petri Nets*, volume 616 of *LNCS*, pages 154–172, 1992.
- [11] M. Feinberg. Lectures on chemical reaction networks. Given at the Math. Research Center, Univ. Wisconsin, 1979. Available online at <http://www.che.eng.ohio-state.edu/~feinberg/LecturesOnReactionNetworks>.
- [12] G. Florin and S. Natkin. Generalization of queueing network product form solutions to stochastic Petri nets. *IEEE Trans. Software Engrg.*, 17(2):99–107, 1991.
- [13] W. Gordon and G. Newell. Closed queueing systems with exponential servers. *Oper. Res.*, 15:254–265, 1967.
- [14] S. Haddad, P. Moreaux, M. Sereno, and M. Silva. Product-form and stochastic Petri nets: a structural approach. *Performance Evaluation*, 59(4):313 – 336, 2005.
- [15] W. Henderson, D. Lucic, and P. Taylor. A net level performance analysis of stochastic Petri nets. *J. Austral. Math. Soc. Ser. B*, 31(2):176–187, 1989.
- [16] J.R. Jackson. Networks of waiting lines. *Oper. Res.*, 5:518–521, 1957.
- [17] F. Kelly. *Reversibility and Stochastic Networks*. Wiley, New-York, 1979.
- [18] T. G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *Journal of Chemical Physics*, 57(7):2976–2978, 1972.
- [19] A. A. Lazar and T. G. Robertazzi. Markovian Petri net protocols with product form solution. *Performance Evaluation*, 12(1):67 – 77, 1991.
- [20] J. Mairesse and H.-T. Nguyen. Deficiency zero Petri nets and product form. In G. Franceschinis and K. Wolf, editors, *Petri Nets 2009*, volume 5606 of *LNCS*, pages 103–122. Springer-Verlag, 2009.
- [21] T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541 – 580, 1989.
- [22] C.A. Petri and W. Reisig. Petri net. *Scholarpedia*, 3(4):6477, 2008.
- [23] C. Reutenauer. *The mathematics of Petri nets*. Prentice Hall, 1990.
- [24] R. Serfozo. *Introduction to stochastic networks*, volume 44 of *Applications of Mathematics*. Springer-Verlag, New York, 1999.