Synthesis and Analysis of Product-form Petri Nets

Serge Haddad
ENS Cachan, LSV,
CNRS UMR 8643, INRIA, Cachan, France
haddad@lsv.ens-cachan.fr

Jean Mairesse, Hoang-Thach Nguyen
Université Paris 7, LIAFA,
CNRS UMR 7089, Paris, France
{mairesse, ngthach}@liafa.jussieu.fr

Abstract

For a large Markovian model, a "product form" is an explicit description of the steady-state behaviour which is otherwise generally untractable. Being first introduced in queueing networks, it has been adapted to Markovian Petri nets. Here we address three relevant issues for product-form Petri nets which were left fully or partially open: (1) we provide a sound and complete set of rules for the synthesis; (2) we characterise the exact complexity of classical problems like reachability; (3) we introduce a new subclass for which the normalising constant (a crucial value for product-form expression) can be efficiently computed.

Keywords: Petri nets, product-form, synthesis, complexity analysis, reachability, normalising constant

1 Introduction

Product-form for stochastic models. Markovian models of discrete events systems are powerful formalisms for modelling and evaluating the performances of such systems. The main goal is the equilibrium performance analysis. It requires to compute the stationary distribution of a continuous time Markov process derived from the model. Unfortunately the potentially huge (sometimes infinite) state space of the models often prevents the modeller from computing explicitly this distribution. To cope with the issue, one can forget about exact solutions and settle for approximations, bounds, or even simulations. The other possibility is to focus on subclasses for which some kind of explicit description is indeed possible. In this direction, the most efficient and satisfactory approach may be the *product-form* method: for a model composed of modules, the stationary probability of a global state may be expressed as a product of quantities depending only on local states divided by a *normalising constant*.

Such a method is applicable when the interactions between the modules are “weak”. This is the case for queueing networks where the interactions between queues are described by a random routing of clients. Various classes of queueing networks with product-form solutions have been exhibited [18, 6, 19]. Moreover efficient algorithms have been designed for the computation of the normalising constant [25].

Product-form Petri nets. Due to the explicit modelling of competition and synchronisation, the Markovian Petri nets formalism [1] is an attractive modelling paradigm. Similarly to queueing networks, product-form Markovian Petri Nets were introduced to cope with the combinatorial explosion of the state space. Historically, works started with purely behavioural properties (i.e. by an analysis of the reachability graph) as in [20], and then progressively moved to more and more structural characterisations [21, 17]. Building on the work of [17], the authors of [14] establish the first purely structural condition for which a product form exists and propose a polynomial time algorithm to check for the condition, see also [22] for an alternative characterisation. These nets are called Π^2-nets.

Product-form Petri nets have been applied for the specification and analysis of complex systems. From a modelling point of view, compositional approaches have been proposed [3, 5] as well as hierarchical ones [16]. Application fields have also been identified like (1) hardware design and more particularly RAID storage [16] and (2) software architectures [4].

Open issues related to product-form Petri nets.

- From a modelling point of view, it is more interesting to design specific types of Petri nets by modular constructions rather than checking a posteriori whether a net satisfies the specification. For instance, in [12], a sound and complete set of rules is proposed for the synthesis of live and bounded free-choice nets. Is it possible to get an analog for product-form Petri nets?
- From a qualitative analysis point of view, it is interesting to know the complexity of classical problems (reachability, coverability, liveness, etc.) for a given subclass of Petri nets and to compare it with that of general Petri nets. For product-form Petri nets, partial results were presented in [14] but several questions were left open. For instance, the reachability problem is PSPACE-complete for safe Petri nets but in safe product-form Petri nets it is only proved to be NP-hard in [14].
- From a quantitative analysis point of view, an important and difficult issue is the computation of the normalising constant. Indeed, in product-form Petri nets, one can directly compute relative probabilities (e.g. available versus unavailable service), but determining absolute probabilities requires to compute the normalising constant (i.e. the sum over reachable states of the relative probabilities). In

models of queueing networks, this can be efficiently performed using dynamic programming. In Petri nets, it has been proved that the efficient computation is possible when the linear invariants characterise the set of reachable markings [11]. Unfortunately, all the known subclasses of product-form nets that fulfill this characterisation are models of queueing networks!

Our contribution. Here we address the three above issues. In Section 3, we provide a set of sound and complete rules for generating any Π^2-net. We also use these rules for transforming a general Petri net into a related product-form Petri net. In Section 4, we solve relevant complexity issues. More precisely, we show that the reachability and liveness problems are PSPACE-complete for safe product-form nets and that the coverability problem is EXPSPACE-complete for general product-form nets. From these complexity results, we conjecture that the problem of computing the normalising constant does not admit an efficient solution for the general class of product-form Petri nets. However, in Section 5, we introduce a large subclass of product-form Petri nets, denoted Π^3-nets, for which the normalising constant can be efficiently computed. We emphasise that contrary to all subclasses related to queueing networks, Π^3-nets may admit *spurious* markings (i.e. that fufill the invariants while being unreachable).

The above results may change our perspective on product-form Petri nets. It is proved in [22] that the intersection of free-choice and product-form Petri nets is the class of Jackson networks [18]. This may suggest that the class of product-form Petri nets is somehow included in the class of product-form queueing networks. In the present paper, we refute this belief in two ways. First by showing that some classical problems are as complex for product-form Petri nets as for general Petri nets whereas they become very simple for product-form queueing networks. Second by exhibiting the large class of Π^3-nets which can model complex behaviours (e.g. illustrated by the presence of spurious markings).

A conference version of the paper appeared in [15]. The present version includes additional results (Subsection 2.2) together with full proofs of the results. (There is one exception, Proposition 4.3, for which the proof can be found in the arXiv version of the paper available at `http://arxiv.org/abs/1104.0291`)

Notations. We often denote a vector $u \in \mathbb{R}^S$ by $\sum_s u(s)s$. The *support* of vector u is the subset $S' \equiv \{s \in S \mid u(s) \neq 0\}$.

2 Petri nets, product-form nets, and Π^2-nets

Definition 2.1 (Petri net). *A* Petri net *is a 5-tuple* $\mathcal{N} = (P, T, W^-, W^+, m_0)$ *where:*

- P *is a finite set of* places*;*

- T *is a finite set of* transitions*, disjoint from* P*;*
- W^-*, resp.* W^+*, is a* $P \times T$ *matrix with coefficients in* $\mathbb{N}$*;*
- $m_0 \in \mathbb{N}^P$ *is the* initial marking.

Below, we also call *Petri net* the unmarked quadruple (P, T, W^-, W^+). The presence or absence of a marking will depend on the context.

A Petri net is represented in Figure 1. The following graphical conventions are used: places are represented by circles and transitions by rectangles. There is an arc from $p \in P$ to $t \in T$ (resp. from $t \in T$ to $p \in P$) if $W^+(p,t) > 0$ (resp. $W^-(p,t) > 0$), and the weight $W^+(p,t)$ (resp. $W^-(p,t)$) is written above the corresponding arc except when it is equal to 1 in which case it is omitted. The initial marking is materialised: if $m_0(p) = k$, then k tokens are drawn inside the circle p. Let $P' \subset P$ and m be a marking then $m(P')$ is defined by $m(P') \equiv \sum_{p \in P'} m(p)$.

The matrix $W = W^+ - W^-$ is the *incidence matrix* of the Petri net. The *input bag* ${}^\bullet t$ (resp. *output bag* $t^\bullet$) of the transition t is the column vector of W^- (resp. W^+) indexed by t. For a place p, we define ${}^\bullet p$ and $p^\bullet$ similarly. A *T-semi-flow* (resp. *S-semi-flow*) is a $\mathbb{Q}$-valued vector v such that $W.v = (0, \dots, 0)$ (resp. $v.W = (0, \dots, 0)$).

A *symmetric* Petri net is a Petri net such that: $\forall t \in T,\ \exists t^- \in T,\quad {}^\bullet t = (t^-)^\bullet, t^\bullet = {}^\bullet t^-$. A *free-choice net* is a Petri net such that: $\forall t, t' \in T$, either ${}^\bullet t \cap {}^\bullet t' = \emptyset$, or ${}^\bullet t = {}^\bullet t'$. A *state machine* is a Petri net such that: $\forall t \in T,\ |{}^\bullet t| = |t^\bullet| = 1$. A *marked graph* is a Petri net such that: $\forall p \in P,\ |{}^\bullet p| = |p^\bullet| = 1$.

Definition 2.2 (Firing rule)**.** *A transition* t *is* enabled *by the marking* m *if* $m \geq {}^\bullet t$ *(denoted by* $m \xrightarrow{t}$*); an enabled transition* t *may* fire *which transforms the marking* m *into* $m - {}^\bullet t + t^\bullet$*, denoted by* $m \xrightarrow{t} m' = m - {}^\bullet t + t^\bullet$*.*

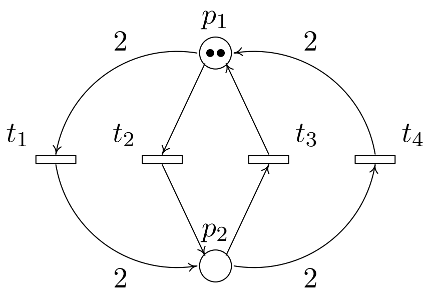

$$W = \begin{pmatrix} -2 & -1 & 1 & 2 \\ 2 & 1 & -1 & -2 \end{pmatrix}.$$
$$m_0 = (2, 0).$$

Figure 1: Petri net.

A marking m' is *reachable* from the marking m if there exists a *firing sequence* $\sigma = t_1 \dots t_k$ $(k \geq 0)$ and a sequence of markings $m_1, \dots, m_{k-1}$ such that $m \xrightarrow{t_1} m_1 \xrightarrow{t_2} \cdots \xrightarrow{t_{k-1}} m_{k-1} \xrightarrow{t_k} m'$. We write in a condensed way: $m \xrightarrow{\sigma} m'$.

We denote by $\mathcal{R}(m)$ the set of markings which are reachable from the marking m. The *reachability graph* of a Petri net with initial marking m_0 is the directed graph with nodes $\mathcal{R}(m_0)$ and arcs $\{(m,m')|\exists t \in T : m \xrightarrow{t} m'\}$.

Given $(\mathcal{N}, m_0)$ and m_1, the *reachability problem* is to decide if $m_1 \in \mathcal{R}(m_0)$, and the *coverability problem* is to decide if $\exists m_2 \in \mathcal{R}(m_0), m_2 \geq m_1$.

A Petri net $(\mathcal{N}, m_0)$ is *live* if every transition can always be enabled again, that is: $\forall m \in \mathcal{R}(m_0), \forall t \in T,\ \exists m' \in \mathcal{R}(m),\ m' \xrightarrow{t}$. A Petri net $(\mathcal{N}, m_0)$ is bounded if $\mathcal{R}(m_0)$ is finite. It is *safe* or *1-bounded* if: $\forall m \in \mathcal{R}(m_0),\ \forall p \in P,\ m(p) \leq 1$.

2.1 Product-form Petri nets

There exist several ways to define timed models of Petri nets, see [2]. We consider the model of Markovian Petri nets with *race policy*. Roughly, with each enabled transition is associated a "countdown clock" whose positive initial value is set at random according to an exponential distribution whose rate depends on the transition. The first transition to reach 0 fires, which may enable new transitions and start new clocks. We adopt here the *single-server policy* which means that the rate of a transition does not depend on the enabling degree of the transition. In the more general definition of product-form Petri nets [14, Definition 8], rates may depend on the current marking in a restricted way. For the sake of readability, we have chosen a simpler version. Results of sections 3 and 4 still hold with the general definition. On the other hand, it is well-known that the complexity of the computation of the normalisation constant highly increases even for the simple case of queuing networks. Here also the results of section 5 are only valid with constant rates.

Definition 2.3 (Markovian PN)**.** *A* Markovian Petri net (with race policy) *is a Petri net equipped with a set of* rates $(\mu_t)_{t\in T}$*,* $\mu_t \in \mathbb{R}_+^*$*. The firing time of an enabled transition* t *is exponentially distributed with parameter* μ_t*. The marking evolves as a continuous-time jump Markov process with state space* $\mathcal{R}(m_0)$ *and infinitesimal generator* $Q = (q_{m,m'})_{m,m' \in \mathcal{R}(m_0)}$*, given by:*

$$\forall m,\ \forall m' \neq m,\ q_{m,m'} = \sum_{t \text{ such that } m \xrightarrow{t} m'} \mu_t, \qquad \forall m,\ q_{m,m} = -\sum_{m' \neq m} q_{m,m'} \,. \tag{2.1}$$

W.l.o.g., we assume that there is no transition t such that ${}^\bullet t = t^\bullet$. Indeed, the firing of such a transition does not modifiy the marking, so its removal does not modify the infinitesimal generator. We also assume that $({}^\bullet t_1, t_1^\bullet) \neq ({}^\bullet t_2, t_2^\bullet)$ for all transitions $t_1 \neq t_2$. Indeed, if it is not the case, the two transitions may be replaced by a single one with the summed rate.

An *invariant measure* is a non-trivial solution ν to the *balance equations*: $\nu Q = (0, \ldots, 0)$. A *stationary distribution* π is an invariant probability measure: $\pi Q = (0, \ldots, 0)$, $\sum_m \pi(m) = 1$.

Definition 2.4 (Product-form PN)**.** *A Petri net is a* product-form Petri net *if for all rates $(\mu_t)_{t \in T}$, the corresponding Markovian Petri net admits an invariant measure ν satisfying:*

$$\exists (u_p)_{p \in P}, u_p \in \mathbb{R}_+, \quad \forall m \in \mathcal{R}(m_0), \qquad \nu(m) = \prod_{p \in P} u_p^{m_p} \, . \tag{2.2}$$

The existence of ν satisfying (2.2) implies that the marking process is irreducible (in other words, the reachability graph is strongly connected). In (2.2), the mass of the measure, i.e. $\nu(\mathcal{R}(m_0)) = \sum_m \nu(m)$, may be either finite or infinite. For a bounded Petri net, the mass is always finite. But for an unbounded Petri net, the typical situation will be as follows: structural conditions on the Petri net will ensure that the Petri net is a product-form one. Then, for some values of the rates, ν will have an infinite mass, and, for others, ν will have a finite mass. In the first situation, the marking process will be either transient or recurrent null (unstable case). In the second situation, the marking process will be positive recurrent (stable or ergodic case).

When the mass is finite, we call $\nu(\mathcal{R}(m_0))$ the *normalising constant*. The probability measure $\pi(\cdot) = \nu(\mathcal{R}(m_0))^{-1} \nu(\cdot)$ is the unique stationary measure of the marking process. Computing explicitly the normalising constant is an important issue, see Section 5.

The goal is now to get sufficient conditions for a Petri net to be of product-form. To that purpose, we introduce three notions: *weak reversibility*, *deficiency*, and *witnesses*.

Let (N, m_0) be a Petri net. The set of *complexes* is defined by $\mathcal{C} = \{{}^\bullet t \mid t \in T\} \cup \{t^\bullet \mid t \in T\}$. The *reaction graph* is the directed graph whose set of nodes is $\mathcal{C}$ and whose set of arcs is $\{({}^\bullet t, t^\bullet) | t \in T\}$. It can be viewed as a state machine.

Definition 2.5 (Weak reversibility: Π-nets)**.** *A Petri net is* weakly reversible (WR) *if every connected component of its reaction graph is strongly connected. Weakly reversible Petri nets are also called* Π-nets*.*

The notion and the name "WR" come from the chemical literature. In the Petri net context, it was introduced in [7, Assumption 3.2] under a different name and with a slightly different but equivalent formulation. WR is a strong constraint. It should not be confused with the classical notion of "reversibility" (the marking graph is strongly connected). In particular WR, a structural property, implies reversibility, a behavioural one! Observe that all symmetric Petri nets are WR.

The notion of deficiency is due to Feinberg [13].

Definition 2.6 (Deficiency)**.** *Consider a Petri net with incidence matrix W and set of complexes $\mathcal{C}$. Let ℓ be the number of connected components of the reaction graph. The* deficiency *of the Petri net is defined by: $|\mathcal{C}| - \ell - rank(W)$.*

The notion of witnesses appears in [14].

Definition 2.7 (Witness)**.** *Let c be a complex. A* witness *of c is a vector $wit(c) \in \mathbb{Q}^P$ such that for all transition t:*

$$\begin{cases} wit(c) \cdot W(t) = -1 & \text{if } {}^\bullet t = c \\ wit(c) \cdot W(t) = 1 & \text{if } t^\bullet = c \\ wit(c) \cdot W(t) = 0 & \text{otherwise}, \end{cases}$$

where $W(t)$ denotes the column vector of W indexed by t.

Examples. Consider the Petri net of Figure 1. First, it is WR. Indeed, the set of complexes is $\mathcal{C} = \{p_1, p_2, 2p_1, 2p_2\}$ and the reaction graph is:

$$p_1 \leftrightarrow p_2\,,\ 2p_1 \leftrightarrow 2p_2\,,$$

with two connected components which are strongly connected. Second, the deficiency is 1 since $|\mathcal{C}| = 4$, $\ell = 2$, and $\text{rank}(W) = 1$. Last, one can check that none of the complexes admit a witness.

The Petri net of Figure 4 is WR and has deficiency 0. Note that the witnesses may not be unique. Possible witnesses are: $wit(2p_1 + q_1) = q_1$, $wit(p_1 + q_2) = q_2$, $wit(p_2 + q_3) = q_3$, $wit(2p_2 + q_4) = q_4$. Another possible set of witnesses is $\{q_1, q_2, -q_2, -q_1\}$.

Proposition 2.8 (deficiency 0 $\iff$ witnesses, in [22, Prop. 3.9])**.** *A Petri net admits a witness for each complex iff it has deficiency 0.*

Next Theorem is a combination of Feinberg's Deficiency zero Theorem [13] and Kelly's Theorem [19, Theorem 8.1]. (It is proved under this form in [22, Theorem 3.8].)

Theorem 2.9 (WR + deficiency 0 $\implies$ product-form)**.** *Consider a Markovian Petri net with rates $(\mu_t)_{t\in T}$, $\mu_t > 0$, and assume that the underlying Petri net is WR and has deficiency 0. Then there exists $(u_p)_{p\in P}$, $u_p > 0$, satisfying the equations:*

$$\forall c \in \mathcal{C}, \qquad \prod_{p:c_p \neq 0} u_p^{c_p} \sum_{t:{}^\bullet t = c} \mu_t = \sum_{t:t^\bullet = c} \mu_t \prod_{p:{}^\bullet t_p \neq 0} u_p^{{}^\bullet t_p}\,. \tag{2.3}$$

The marking process has an invariant measure ν such that:

$$\forall m,\ \nu(m) = \Phi(m)^{-1} \prod_{p\in P} u_p^{m_p}\,.$$

Checking the WR, computing the deficiency, determining the witnesses, and solving the equations (2.3), all of these operations can be performed in polynomial-time, see [14, 22].

Summing up the above, it seems worth to isolate and christen the class of nets which are WR and have deficiency 0. We adopt the terminology of [14].

Definition 2.10 (Π^2-net)**.** *A* Π^2-net *is a Petri net which is WR and has deficiency 0.*

2.2 Some properties of WR and deficiency zero nets

Let $\mathcal{N} = (P, T, W^-, W^+)$ be a Petri net. Let $W = W^+ - W^-$ be the incidence matrix of $\mathcal{N}$ and let A be the incidence matrix of the reaction graph.

Consider at first free-choice nets. It was shown in [22, Section 4.3] that for free-choice nets, WR implies deficiency zero. The converse does not hold for general free-choice nets. For instance, state machines always have deficiency zero [22, Prop. 3.2], and may not be WR. For marked graphs, however, the converse is true, and stated below.

Proposition 2.11. *The deficiency of a connected marked graph is either 0 or 1. A marked graph has deficiency zero if and only if it is WR.*

Proof. Let $\mathcal{N}$ be a marked graph. According to [9, Prop. 3.16], the only T-semi-flows of $\mathcal{N}$ are $a(1, \cdots, 1)$, $a \in \mathbb{Q}$, hence $\text{rank}(W) = |T| - 1$. Since A is a $\mathcal{C} \times T$ matrix, $\text{rank}(A) \leq |T|$. Hence $\delta = \text{rank}(A) - \text{rank}(W) \leq 1$.

The "if" direction of the second claim is trivial since a marked graph is a free-choice net. Consider the "only if" direction. Let $\mathcal{N}$ be a deficiency zero marked graph. Let $\mathbf{1}$ be the column vector $(1, \ldots, 1)$ of size T. Since $\mathcal{N}$ is a marked graph, we have $W \cdot \mathbf{1} = (0, \ldots, 0)$. By Proposition 2.8, $A = BW$ for some $\mathbb{Q}$-valued matrix B. So we have $A \cdot \mathbf{1} = BW \cdot \mathbf{1} = (0, \ldots, 0)$. This implies that the connected components of the reaction graph must be strongly connected. Indeed pick a connected component which is not strongly connected. It admits a partition of its complexes into two subsets C_1 and C_2 such that there at least one transition t from C_1 to C_2 and no transition from C_2 to C_1. Then vector x defined by $x(c) = 0$ for $c \in C_1$ and $x(c) = 1$ for $c \in C_2$ fulfills $x.A \geq 0$ and $x.A(t) > 0$. Thus $x.A.\mathbf{1} > 0$ yields a contradiction. So $\mathcal{N}$ is WR. ☐

Proposition 2.12. *For a live and bounded Petri net, deficiency zero implies weak reversibility.*

Proof. Let m_o be a marking such that $(\mathcal{N}, m_0)$ is live and bounded. We assume that $\mathcal{N}$ has deficiency 0 but is not WR. Then there exists a terminal

strongly connected component C of the reaction graph and a transition t_0 such that $t_0^\bullet \in C$ and ${}^\bullet t_0 \notin C$.
We claim that for every vector $v \in \mathbb{Q}^T$ such that for all $t \in T$, $v(t) \geq 0$ and $v(t_0) > 0$, we have $Av \neq (0,\ldots,0)$. Indeed,

$$\begin{aligned}\sum_{c\in C}(Av)(c) &= \sum_{c\in C}\left(\sum_{t\in T} v(t)\big(\mathbf{1}_{t^\bullet=c} - \mathbf{1}_{{}^\bullet t=c}\big)\right)\\ &= \sum_{t\in T-\{t_0\}} v(t)\left(\sum_{c\in C}\big(\mathbf{1}_{t^\bullet=c} - \mathbf{1}_{{}^\bullet t=c}\big)\right) + v(t_0)\,.\end{aligned}$$

Since C is a terminal strongly connected component, $\sum_{c\in C} \mathbf{1}_{t^\bullet=c} - \mathbf{1}_{{}^\bullet t=c}$ is either 0 or 1 for all $t \in T$. Hence $\sum_{c\in C}(Av)(c) \geq v(t_0) > 0$. The claim is proved.

Since $(\mathcal{N}, m_0)$ is live and bounded, there exists a strictly positive T-semiflow $v \in \mathbb{Q}^T$ [9, Theorem 2.38], that is: $\forall t,\ v(t) > 0,\ W \cdot v = (0,\ldots,0)$. Now recall that the deficiency of $\mathcal{N}$ is 0. According to Proposition 2.8, there exists a $\mathcal{C}\times P$ matrix B such that $A = BW$. We get $Av = BWv = (0,\ldots,0)$. This contradicts the above claim. □

A *home marking* is a marking which is reachable from every reachable marking. Having a home marking is an important property for Markovian Petri nets. Indeed, a Petri net has a home marking iff its reachability graph has only one terminal strongly connected component. And this last condition is required for the marking process to be ergodic.

Proposition 2.13. *Let $\mathcal{N}$ be a deficiency zero Petri net. Then $\mathcal{N}$ is WR iff there exists a marking m_0 such that $(\mathcal{N}, m_0)$ is live and m_0 is a home marking.*

Proof. Suppose that $\mathcal{N}$ is WR. Let m_0 be a marking which enables every transition. The definition of weak reversibility implies that every arc of the reachability graph belongs to a cycle, so the reachability graph is strongly connected, that is m_0 is a home marking. The liveness follows trivially.

Now suppose that there exists a marking m_0 such that $(\mathcal{N}, m_0)$ is live and m_0 is a home marking but $\mathcal{N}$ is not WR. We proceed as in the proof of Prop. 2.12. Let C be a terminal strongly connected component of the reaction graph and let t be a transition such that $t^\bullet \in C$ and ${}^\bullet t \notin C$. Since $(\mathcal{N}, m_0)$ is live there is a path γ_1 in the reachability graph from m_0 to m_1 which enables t. Let m_1' be the marking reached by the firing of t, since m_0 is a home marking there is a path γ_2 from m_1' to m_0. Thus $\gamma = \gamma_1 t \gamma_2$ is a (directed) cycle of the reachability graph of $(\mathcal{N}, m_0)$. Let v be the $\mathbb{N}^T$ column vector such that: $\forall u \in T$, $v(u)$ is the number of occurrences of u in γ. Clearly, $v(t) > 0$ and $W.v = (0,\ldots,0)$. The end of the argument follows from the claim inside the proof of Prop. 2.12. □

The interest of Prop. 2.13 is twofold. On the one hand, it connects weak reversibility and deficiency zero which are two independent properties ([22]). On the other hand, it shows that the only deficiency zero and live Markovian Petri nets which are ergodic are the Π^2-nets.

Figure 2 recapitulates the relations between deficiency and weak reversibility. The shaded cells correspond to impossibilities. For instance, no WR free-choice nets have strictly positive deficiency.

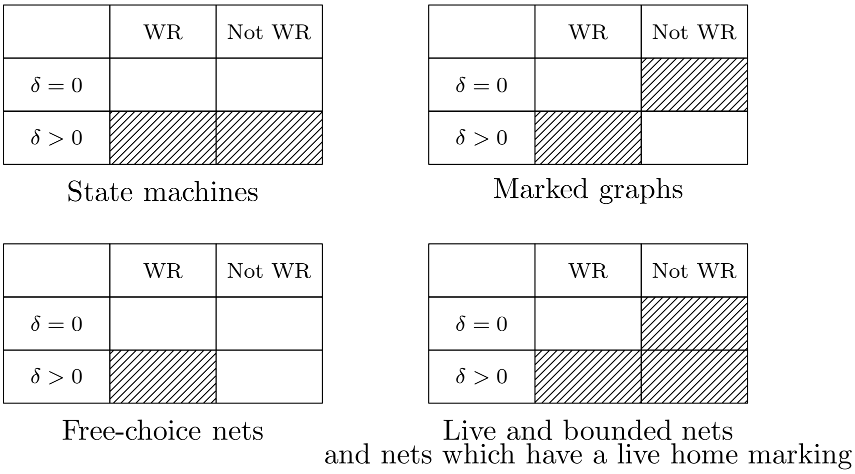

Figure 2: Relations between deficiency (δ) and WR for some classes of Petri nets.

3 Synthesis and regulation of Π^2-nets

The reaction graph, defined in Section 2.1, may be viewed as a Petri net (state machine). Let us formalise this observation. The *reaction Petri net* of $\mathcal{N}$ is the Petri net $\mathcal{A} = (\mathcal{C}, T, \overline{W}^-, \overline{W}^+)$, with for every $t \in T$:

- $\overline{W}^-({}^\bullet t, t) = 1$ and $\forall u \neq {}^\bullet t,\ \overline{W}^-(u, t) = 0$
- $\overline{W}^+(t^\bullet, t) = 1$ and $\forall u \neq t^\bullet,\ \overline{W}^+(u, t) = 0$

3.1 Synthesis

In this subsection, we consider unmarked nets. We define three rules that generate all the Π^2-nets. The first rule adds a strongly connected state machine.

Definition 3.1 (State-machine insertion). *Let $\mathcal{N} = (P_\mathcal{N}, T_\mathcal{N}, W^-_\mathcal{N}, W^+_\mathcal{N})$ be a net and $\mathcal{M} = (P_\mathcal{M}, T_\mathcal{M}, W^-_\mathcal{M}, W^+_\mathcal{M})$ be a strongly connected state machine disjoint from $\mathcal{N}$. The rule* `S-add` *is always applicable and $\mathcal{N}' =$* `S-add`$(\mathcal{N}, \mathcal{M})$ *is defined by:*

- $P' = P_{\mathcal{N}} \sqcup P_{\mathcal{M}}$, $T' = T_{\mathcal{N}} \sqcup T_{\mathcal{M}}$;
- $\forall p \in P_{\mathcal{N}},\ \forall t \in T_{\mathcal{N}},\ W'^{-}(p,t) = W^{-}_{\mathcal{N}}(p,t),\ W'^{+}(p,t) = W^{+}_{\mathcal{N}}(p,t)$;
- $\forall p \in P_{\mathcal{M}},\ \forall t \in T_{\mathcal{M}},\ W'^{-}(p,t) = W^{-}_{\mathcal{M}}(p,t),\ W'^{+}(p,t) = W^{+}_{\mathcal{M}}(p,t)$;
- *All other entries of W'^{-} and W'^{+} are null.*

The second rule consists in substituting to a complex c the complex $c+\lambda p$. However in order to be applicable some conditions must be fulfilled. The first condition requires that $c(p)+\lambda$ is non-negative. The second condition ensures that the substitution does not modify the reaction graph. The third condition preserves deficiency zero. Observe that the third condition can be checked in polynomial time, indeed it amounts to solving a system of linear equations in $\mathbb{Q}$ for every complex.

Definition 3.2 (Complex update)**.** *Let $\mathcal{N} = (P, T, W^{-}, W^{+})$ be a Π^2-net, c be a complex of $\mathcal{N}$, $p \in P$, $\lambda \in \mathbb{Z} \setminus \{0\}$. The rule* `C-update` *is applicable when:*

1. $\lambda + c(p) \geq 0$;
2. $c + \lambda p$ *is not a complex of* $\mathcal{N}$;
3. *For every complex c' there exists a witness $wit(c')$ s.t. $wit(c')(p) = 0$.*

The resulting net $\mathcal{N}' =$ `C-update`$(\mathcal{N}, c, p, \lambda)$ is defined by:

- $P' = P$, $T' = T$;
- $\forall t \in T$ *s.t.* $W^{-}(t) \neq c,\ W'^{-}(t) = W^{-}(t)$, $\forall t \in T$ *s.t.* $W^{-}(t) = c,\ W'^{-}(t) = c + \lambda p$
- $\forall t \in T$ *s.t.* $W^{+}(t) \neq c,\ W'^{+}(t) = W^{-}(t)$, $\forall t \in T$ *s.t.* $W^{+}(t) = c,\ W'^{+}(t) = c + \lambda p$.

The last rule "cleans" the net by deleting an isolated place. We call this operation `P-delete`.

Definition 3.3 (Place deletion)**.** *Let $\mathcal{N} = (P, T, W^{-}, W^{+})$ be a net and let p be an isolated place of $\mathcal{N}$, i.e. $W^{-}(p) = W^{+}(p) = 0$. Then the rule* `P-delete` *is applicable and $\mathcal{N}' =$* `P-delete`$(\mathcal{N}, p)$ *is defined by:*

- $P' = P \setminus \{p\}$, $T' = T$;
- $\forall q \in P',\ W'^{-}(q) = W^{-}(q),\ W'^{+}(q) = W^{+}(q)$.

Proposition 3.4 shows the interest of the rules for synthesis of Π^2-nets.

Proposition 3.4 (Soundness and Completeness)**.** *Let $\mathcal{N}$ be a Π^2-net.*

- *If a rule* `S-add`, `C-update` *or* `P-delete` *is applicable on $\mathcal{N}$ then the resulting net is still a Π^2-net.*
- *The net $\mathcal{N}$ can be obtained by successive applications of the rules* `S-add`, `C-update`, `P-delete` *starting from the empty net.*

Proof. **Soundness.** The case of `P-delete` is straightforward. Since we delete an isolated place, the reaction graph is unchanged. So the net is still WR. Assume that we delete an isolated place p and that p occurs in a witness $wit(c)$ of some complex c. Then $wit(c) - wit(c)(p)$ is also a witness of c.

Let us examine the application of rule $\texttt{S-add}(\mathcal{N}, \mathcal{M})$. The state machine $\mathcal{M}$ constitutes a new component of the reaction graph. Since $\mathcal{M}$ is strongly connected, the new net is still WR. The witness of complexes associated with $\mathcal{N}$ are unchanged. Let q be a place of $\mathcal{M}$; by definition of state machines this place is self-witnessing i.e. $wit(q) = q$. Thus the new net has deficiency zero.

Let us examine the application of the rule $\texttt{C-update}(\mathcal{N}, c, p, \lambda)$. By the second condition of its application the reaction graph of the new net is the same as the original one (with $c + \lambda p$ instead of c). So the new net is WR. Due to the third condition, the witness of $c' \neq c$ is unchanged and the witness of $c + \lambda \cdot p$ is the one of c.

Completeness. Let $\mathcal{N} = (P, T, W^-, W^+)$ be a Π^2-net. We proceed as follows to generate $\mathcal{N}$ via our rules. At any stage of the generation, $\mathcal{N}_{cur}$ denotes the current net. Initially $\mathcal{N}_{cur}$ is the empty net.

First step. Let $\mathcal{A}_1, \ldots, \mathcal{A}_n$ be the strongly connected state machines corresponding to the components of the reaction net of $\mathcal{N}$. Given a complex c of $\mathcal{N}$, the corresponding place in the state machine is denoted q_c. We apply the rules $\texttt{S-add}(\mathcal{N}_{cur}, \mathcal{A}_i)$ for i from 1 to n. At this stage, $\mathcal{N}_{cur}$ has T for set of transitions and a place q_c for every complex c of $\mathcal{N}$. Furthermore, q_c has for input (resp. output) transitions the input (resp. output) transitions of c in $\mathcal{N}$. The complexes of $\mathcal{N}_{cur}$ are the places q_c and they are their own witnesses.

Second step. It consists in adding the places of P in such a way that the net $\mathcal{N}_{cur}$ restricted to the places of P is $\mathcal{N}$. At every stage of this step, given a complex $c = \sum_{p \in P} c(p)p$ of $\mathcal{N}$, there is a corresponding complex $c' = q_c + \sum_{p \in P \cap P_{cur}} c(p)p$ in $\mathcal{N}_{cur}$. For every place $p \in P$, we add p to $\mathcal{N}_{cur}$ by rule $\texttt{S-add}$ (an isolated place is a strongly connected state machine) and for every complex c of $\mathcal{N}$ such that $c(p) > 0$, we apply the rule $\texttt{C-update}(\mathcal{N}_{cur}, c', p, c(p))$. Let us check that this rule is applicable. First, $c'(p) + c(p) = c(p)$ is positive. Second, $c' + c(p)p$ is not a complex of $\mathcal{N}_{cur}$ by construction. Third, for every complex c' of $\mathcal{N}_{cur}$, there is a witness consisting in the single place q_c which is in a state machine $\mathcal{A}_i$ (thus different from p). At the end of this step, $\mathcal{N}_{cur}$ is the net $\mathcal{N}$ enlarged with the places of the state machines $\mathcal{A}_i$. Otherwise stated, every complex c' of $\mathcal{N}_{cur}$ is equal to $c + q_c$.

Third step. This step consists in deleting the places of the state machines. We observe that the place q_c only occurs in the complex $c + q_c$. The net $\mathcal{N}$ being a Π^2-net, every complex c' has a witness $wit(c')$ in $\mathcal{N}$. Then $wit(c')$ is a witness for $c' + q_{c'}$ in $\mathcal{N}_{cur}$ whose support does not contain q_c. Thus the rule $\texttt{C-update}(\mathcal{N}_{cur}, c + q_c, q_c, -1)$ is applicable. After its application,

q_c becomes isolated and can be deleted by the rule $\texttt{P-delete}(\mathcal{N}_{cur}, q_c)$. At the end, we have obtained $\mathcal{N}$. □

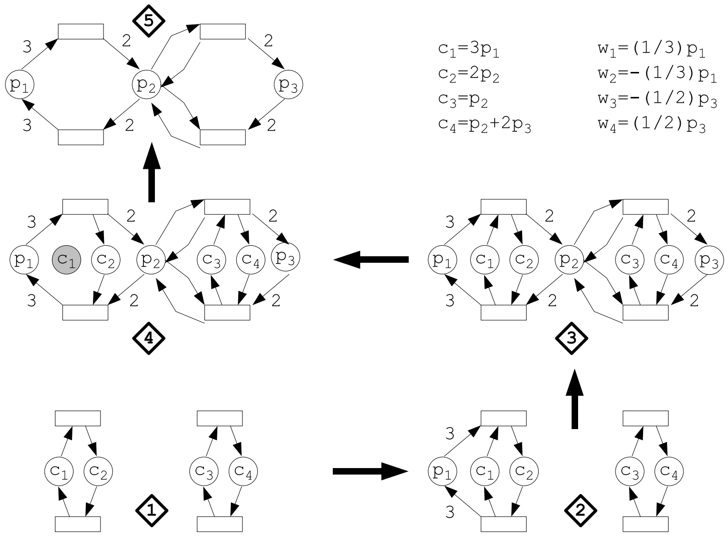

Figure 3: How to synthetise a Π^2-net.

Example. We illustrate the synthesis process using our rules on the net numbered 5 in Figure 3. We have also indicated on the right upper part of this figure, the four complexes and their witnesses. Since the reaction Petri graph of this net has two state machines, we start by creating it using twice the insertion of a state machine (net 1). Then we add the place p_1 (a particular state machine). We update the complex c_1 (the single one where p_1 appears in the original net) by adding $3p_1$ (net 2). Iterating this process, we obtain the net 3. Observe that this net is a fusion (via T the set of transitions) of the original net and its reaction Petri net. We now iteratively update the complexes. The net 4 is the result of transforming $c_1 + 3p_1$ into $3p_1$. Once c_1 is isolated, we delete it. Iterating this process yields the original net.

For modelling purposes, we could define more general rules like the refinement of a place by a strongly connected state machine. Here the goal was to design a minimal set of rules.

3.2 From non Π^2-nets to Π^2-nets

Below we propose a procedure which takes as input any Petri net and returns a Π^2-net. The important disclaimer is that the resulting net, although related to the original one, has a different structural and timed behaviour.

So it is up to the modeller to decide if the resulting net satisfies the desired specifications. In case of a positive answer, the clear gain is that all the associated Markovian Petri nets have a product form.

Consider a Petri net $\mathcal{N} = (P, T, W^-, W^+, m_0)$ with set of complexes $\mathcal{C}$. Assume that $\mathcal{N}$ is not WR. For each transition t, add a reverse transition t^- such that ${}^\bullet t^- = t^\bullet$ and $(t^-)^\bullet = {}^\bullet t$ (unless such a transition already exists). The resulting net is WR. In the Markovian Petri net, the added reverse transitions can be given very small rates, to approximate more closely the original net. However, there is no theoretical guarantee of the convergence of steady-state distributions and in fact counter-examples can be exhibited.

Now, to enforce deficiency 0, the idea is to compose a general Petri net with its reaction graph as in the illustration of Proposition 3.4.

Definition 3.5. *Consider a Petri net $\mathcal{N} = (P, T, W^-, W^+, m_0)$. Let $\overline{m}_0$ be an initial marking for the reaction Petri net $\mathcal{A}$. The* regulated *Petri net associated with $\mathcal{N}$ is defined as follows:*

$$\mathcal{A} \odot \mathcal{N} = \big(P \sqcup \mathcal{C}, T, \widetilde{W}^-, \widetilde{W}^+, (m_0, \overline{m}_0)\big), \quad \widetilde{W}^- = \begin{bmatrix} W^- \\ \overline{W}^- \end{bmatrix}, \widetilde{W}^+ = \begin{bmatrix} W^+ \\ \overline{W}^+ \end{bmatrix}.$$

Proposition 3.6. *The regulated Petri net $\mathcal{A} \odot \mathcal{N}$ is WR iff $\mathcal{N}$ is WR. The regulated Petri net $\mathcal{A} \odot \mathcal{N}$ has deficiency 0.*

Proof. By construction the reaction graph of the regulated Petri net $\mathcal{A} \odot \mathcal{N}$ is the reaction graph of $\mathcal{N}$, i.e. $\mathcal{A}$, modulo a node renaming. So $\mathcal{A} \odot \mathcal{N}$ is WR iff $\mathcal{N}$ is WR.
Now let us prove that the deficiency is 0. We use the characterization by witnesses, see Prop. 2.8. Let $\widetilde{\mathcal{C}}$ be the set of complexes of $\mathcal{A} \odot \mathcal{N}$. Consider $\tilde{c} \in \widetilde{\mathcal{C}}$ and let c be the corresponding element in $\mathcal{C}$. Define $wit(\tilde{c}) \in \mathbb{Q}^{P \sqcup \mathcal{C}}$ by: $wit(\tilde{c})_c = 1, \forall u \neq c, wit(\tilde{c})_u = 0$. By direct inspection, we check that $wit(\tilde{c})$ is indeed a witness of $\tilde{c}$. ☐

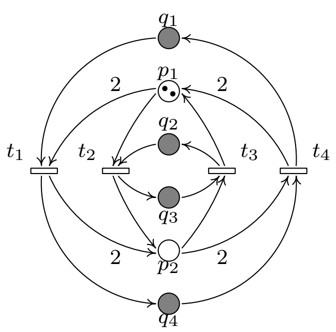

Figure 4: Regulated Petri net associated with the Petri net of Fig 1.

The behaviours of the original and regulated Petri nets are different. In particular, the regulated Petri net is bounded, even if the original Petri net is unbounded. Roughly, the regulation imposes some control on the firing sequences. Consider the example of Figures 1 (original net) and 4 (regulated net). The places q_1, q_2, q_3, q_4 correspond to the complexes $2p_1, p_1, p_2, 2p_2$, respectively. The transitions t_1 and t_4 belong to the same simple circuit in the reaction graph. Let w be an arbitrary firing sequence. The quantity $|w|_{t1} - |w|_{t4}$ is unbounded for the original net, and bounded for the regulated net.

4 Complexity analysis of Π^2-nets

All the nets that we build in this section are symmetric hence WR. For every depicted transition t, the reverse transition exists (sometimes implicitly) and is denoted t^-. It is well known that reachability and liveness of safe Petri nets are PSPACE-complete [10]. In [14], it is proved that reachability and liveness are PSPACE-hard for safe Π-nets and NP-hard for safe Π^2-nets. The next theorem and its corollary improve on these results by showing that the problem is not easier for safe Π^2-nets than for general safe Petri nets.

Theorem 4.1. *The reachability problem for safe Π^2-nets is PSPACE-complete.*

Proof. Our proof of PSPACE-hardness is based on a reduction from the QSAT problem [24]. QSAT consists in deciding whether the following formula is true

$$\varphi \equiv \forall x_n \exists y_n \forall x_{n-1} \exists y_{n-1} \dots \forall x_1 \exists y_1 \psi$$

where ψ is a propositional formula over $\{x_1, y_1 \dots, x_n, y_n\}$ in conjunctive normal form with at most three literals per clause.

Observe that in order to check the truth of φ, one must check the truth of ψ w.r.t. the 2^n interpretations of $x_1, \dots, x_n$ while the corresponding interpretation of any y_i must only depend on the interpretation of $\{x_n, \dots, x_i\}$.

Counters modelling. First we design a Π^2-net $\mathcal{N}_{cnt}$ that "counts" from 0 to $2^k - 1$. This net is defined by:

- $P = \{p_0, \dots, p_{k-1}, q_0, \dots, q_{k-1}\}$;
- $T = \{t_0, \dots, t_{k-1}\}$;
- For every $0 \le i < k$, ${}^\bullet t_i = p_i + \sum_{j<i} q_j$ and $t_i^\bullet = q_i + \sum_{j<i} p_j$;
- For every $0 \le i < k$, $m_0(p_i) = 1$ and $m_0(q_i) = 0$.

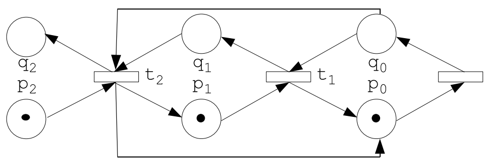

Figure 5: A 3-bit counter (without the reverse transitions).

Observe that for every reachable marking m and every index i, we have $m(p_i) + m(q_i) = 1$. Therefore m can be coded by the binary word $\omega = \omega_{k-1} \dots \omega_0$ in which $\omega_i = m(q_i)$. The word ω is interpreted as the binary expansion of an integer between 0 and $2^k - 1$. We denote by $val(\omega)$ the integer value associated with w. Consider $w \notin \{0^k, 1^k\}$, there are two markings reachable from w which are $w+$ and $w-$ such that $val(w-) = val(w) - 1$ and $val(w+) = val(w) + 1$.

The figure below represents the reachability graph of the 3-bit counter. For a k-bit counter, the shortest firing sequence from 0^k to 1^k is σ_k defined inductively by: $\sigma_1 = t_0$ and $\sigma_{i+1} = \sigma_i t_i \sigma_i$.

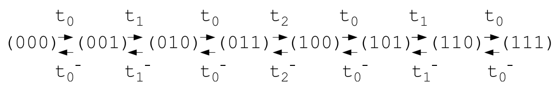

For every complex $c \equiv p_i + \sum_{j<i} q_j$ (resp. $c \equiv q_i + \sum_{j<i} p_j$), a possible witness is $wit(c) \equiv p_i + \sum_{j>i} 2^{j-i-1} p_j$ (resp. $wit(c) \equiv q_i + \sum_{j>i} 2^{j-i-1} q_j$). Thus this subnet has deficiency 0.

To manage transition firings between the update of counters, we duplicate the counter subnet and we synchronize the two subnets as indicated in the figure below. For a duplicated k-bit counter, the shortest firing sequence from the marking with the two counters set to 0^k and place go marked to the marking with the two counters set to 1^k and place go marked is obtained by: $\overline{\sigma}_1 = \overline{t}_0$ and $\overline{\sigma}_{n+1} = \overline{\sigma}_n \overline{t}_n \overline{\sigma}_n$ where $\overline{t}_i = t_i t'_i$.

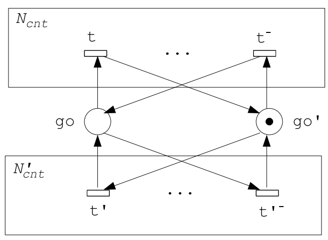

This net has still deficiency 0 since the complexes are just enlarged by the places go or go' and their witnesses remain the same.

Variable modelling. For reasons that will become clear later on, the two counter subnets contain $n+3$ bits indexed from 0 to $n+2$. The bits $1, \dots, n$ of counter *cnt* correspond to the value of variables $x_1, \dots, x_n$. Managing the value of variables $y_1, \dots, y_n$ is done as follows. For every variable y_i, we add the subnet described below on the left (observe that $s_i = r_i^-$) and modify the two counter subnets as described on the right.

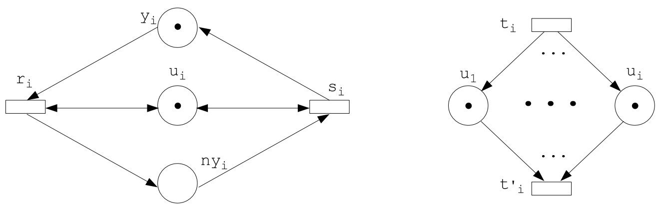

When place y_i (resp. ny_i) is marked, this corresponds to interpreting variable y_i as **true** (resp. **false**). Changes of the interpretation are possible when place u_i is marked. This is the role of the modification done on the counter subnet: between a firing of t_i and t'_i places $\{u_j\}_{j\leq i}$ are marked. With this construction, we get the expected behaviour: the interpretation of a variable y_i can only be modified when the interpretation of a variable x_j with $j \geq i$ is modified. The complexes of the counter subnet are enlarged with places u_i and their witnesses remain the same since places in the support of these witnesses are not modified by transitions s_i and r_i. The new complex $y_i + u_i$ (resp. $ny_i + u_i$) has for witness y_i (resp. ny_i). Thus the new net has still deficiency 0.

Modelling the checking of the propositional formula. We now describe the subnet associated with the checking of propositional formula $\psi \equiv \bigwedge_{j\leq m} C_j$ where we assume w.l.o.g.: (1) that every clause $C_j \equiv l_{j,1} \vee l_{j,2} \vee l_{j,3}$ has exactly three literals (i.e. variables or negated variables); and (2) that every variable or negated variable occurs at least in one clause. The left upper part of Figure 6 shows the Petri net which describes clause C_j of the formula ψ. Places $\ell_{j,k}(k = 1, 2, 3)$ represent the literals while places $n\ell_{j,k}$ represent the literal *used as a proof of the clause*, the place $mutex_j$ avoids to choose several proofs of the clause (and thus ensuring safeness), and finally place $success_j$ can be marked if and only if the evaluation of the clause yields true for the current interpretation and one of its true literal is used as a proof.

The complexes of this subnet are $mutex_j + \ell_{j,k}$ (resp. $success_j + n\ell_{j,k}$) with witness $-n\ell_{j,k}$ (resp. $n\ell_{j,k}$). So the subnet has deficiency 0.

We now synchronise the clause subnets with the previous subnet in order to obtain the final net. Observe that in the previous subnet, transition t_0 (and t'_0) must occur after every interpretation change. This is in fact the role of bit 0 of the counter. Thus we constrain its firing by requiring the places $success_j$ to be marked as presented in the right upper part of Figure 6. Adding loops simply enlarges the complexes associated with t_0 and does not modify the incidence matrix. So the net has still deficiency 0.

It remains to synchronise the value of the variables and the values of the literals where the variables occur either positively or negatively. This is done in two steps. First $\ell_{j,k}$ is initially marked if the interpretation of the initial

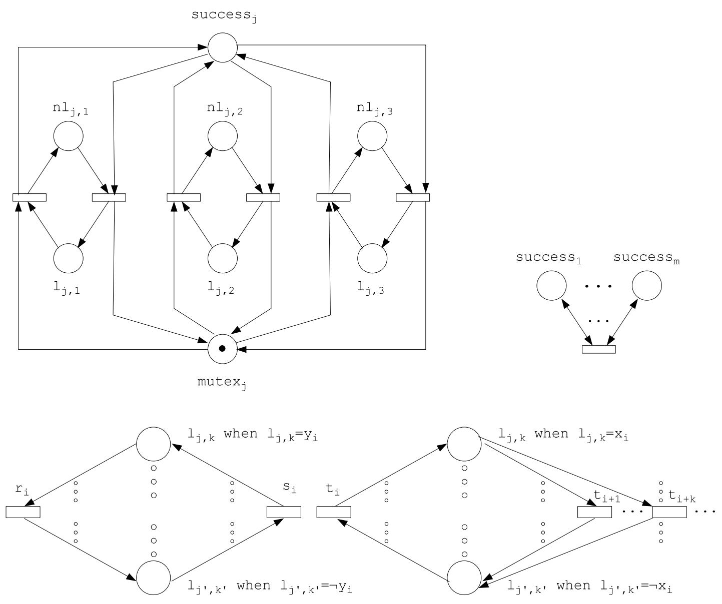

Figure 6: Clause C_j (left), synchronisation with t_0 (right) and with variables (below)

marking satisfies $\ell_{j,k}$. Then we synchronize the value changes as illustrated in the lower part of Figure 6. Once again the complexes are enlarged and the witnesses are still valid since the places $\ell_{j,k}$ do not belong to the support of any witness.

Choice of the initial and final marking for the net. Let us develop a bit the sequence $\overline{\sigma}_{n+3}$ in the two counter subnet in order to explain the choice of initial marking for this subnet:

$$\overline{\sigma}_{n+3} = \overline{\sigma}_{n+1} t_{n+1} t'_{n+1} \overline{\sigma}_{n+1} t_{n+2} t'_{n+2} \overline{\sigma}_{n+1} t_{n+1} t'_{n+1} \overline{\sigma}_{n+1}$$

We want to check all the interpretations of x_i's guessing the appropriate values of y_i's (if they exist). We have already seen that changing from one interpretation to another one (i.e. a counter incrementation or decrementation) allows to perform the allowed updates of y_i. However given the initial interpretation of the x_i's we need to make an initial guess of all the y_i's. So our initial marking restricted to the counter subnet will correspond to the marking reached after $\overline{\sigma}_{n+1} t_{n+1}$, i.e. corresponding to $cnt = 2^{n+1}$ (i.e. word $010\ldots0$), $cnt' = 2^{n+1} - 1$ (i.e. word $001\ldots1$) with in addition places go', u_i's, $mutex_j$'s and y_i's 1-marked; places $\ell_{j,k}$ are marked according to the initial marking of places x_i's and y_i's as explained before. All the other places are unmarked. This explains the role of bit $n+1$.

Furthermore, if we have successfully checked all the interpretations of the x_i's, the counters will have reached the value $2^{n+2}-1$ (corresponding to a firing sequence obtained from $t'_{n+1}\overline{\sigma}_{n+1}$ with possible updates of y_i during change of interpretations). However we do not know what is the final guess for the y_i's. So firing transition t_{n+2} allows to set the y_i's in such a way that the final marking will correspond to $cnt = 2^{n+2}$ (i.e. word $10\dots0$), $cnt' = 2^{n+2}-1$ (i.e. word $01\dots1$) with in addition places go', u_i's $mutex_j$'s and y_i's 1-marked; places $\ell_{j,k}$ are marked accordingly. All the other places are unmarked. This explains the role of bit $n+2$.

By construction, the net reaches the final marking iff the formula is satisfied. Observe that the checking of clauses can be partially done concurrently with the change of interpretation. However as long as, in the net, a clause C_j is "certified" by a literal $\ell_{j,k}$ (i.e. marking place $success_j$ and unmarking place $\ell_{j,k}$) the value of the variable associated with the literal cannot change, ensuring that when t_0 is fired, the marking of any place $success_j$ corresponds to the evaluation of clause C_j with the current interpretation. □

Corollary 4.2. *The liveness problem for safe* Π^2*-nets is* PSPACE*-complete.*

Proof. Observe that the transitions of the net of the previous proof are fireable at least once and so live by reversibility, implied by weak reversibility iff φ is true. □

Let us now consider general (non-safe) Petri nets. Reachability and coverability of symmetric nets is EXPSPACE-complete [23]. In [14], it is proved that both problems are EXPSPACE-complete for WR nets (which include symmetric Petri nets). The next proposition establishes the same result for the coverability of Π^2-nets.

Proposition 4.3. *The coverability problem for* Π^2*-nets is* EXPSPACE*-complete.*

Proof. Since we already know that coverability for Π-nets belongs to EXPSPACE [14], it remains to prove that coverability for Π^2-nets is EXPSPACE-hard. In order to establish this result, we slightly adapt the reduction given in [23] of the termination problem for a three counter machine where the values of counters are bounded by $e_n \equiv 2^{2^n}$ with n the size of (a representation of) the machine. Thus we first depict the original reduction and then we describe our modifications and explain why the reduction is still valid.

For a uniform presentation of the proof we assume w.l.o.g. that the machine has four counters (these more powerful machines include the original ones). The key ingredient is the concise management of counters and more precisely the zero test. Indeed one models a counter c_i with $i \in \{1,2,3,4\}$ by two complementary places $A_{i,n}$ and $B_{i,n}$. When the counter has value x, place $A_{i,n}$ contains x tokens and place $B_{i,n}$ contains $e_n - x$ tokens. Testing (and decrementing) that the counter c_i is greater than 0 is done as usual by an arc with weight 1 starting from $A_{i,n}$. However this approach does not work for the zero test as it would require a (double) arc from $B_{i,n}$ with weight e_n thus implying a net representation of size at least 2^n which would not be valid.

Thus the zero test is managed by an inductive construction (w.r.t. n) of "nested" subnets $\mathcal{N}_k$ leading to a subnet (the union of these subnets) with size in $O(n)$. Let us describe this construction. The main places are: $B_{i,k}$ with $i \in \{1,2,3,4\}, 0 \leq k \leq n$ containing at most e_k tokens and safe places $C_{i,k}$, F_k and S_k. The inductive properties are the following ones:

- In subnet $\bigcup_{l \leq k} \mathcal{N}_l$, starting from marking $S_k + C_{i,k}$ one may reach marking $F_k + C_{i,k} + e_k B_{i,k}$.
- Furthermore any marking reachable from $S_k+C_{i,k}+\alpha_{i_1}B_{i_1,k}+\alpha_{i_2}B_{i_2,k}+\alpha_{i_3}B_{i_3,k}$ ($\{i_1,i_2,i_3\} = \{1,2,3,4\} \setminus \{i\}$) with S_k or F_k marked is either $S_k+C_{i,k}+\alpha_{i_1}B_{i_1,k}+\alpha_{i_2}B_{i_2,k}+\alpha_{i_3}B_{i_3,k}$ or $F_k+C_{i,k}+e_kB_{i,k}+\alpha_{i_1}B_{i_1,k}+\alpha_{i_2}B_{i_2,k}+\alpha_{i_3}B_{i_3,k}$.

Basic case $k = 0$. This case is straightforward: $\mathcal{N}_0$ consists in four transitions when transition corresponding to i is figured below.

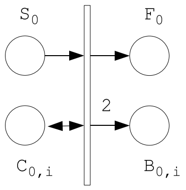

Inductive case. Assume that the inductive properties holds for k. The net corresponding to $\mathcal{N}_{k+1}$ is described below with the following convention: S corresponds to S_{k+1} and s corresponds to S_k. The same convention applies to all names. Furthermore for sake of readability we have duplicated some places in the figure.

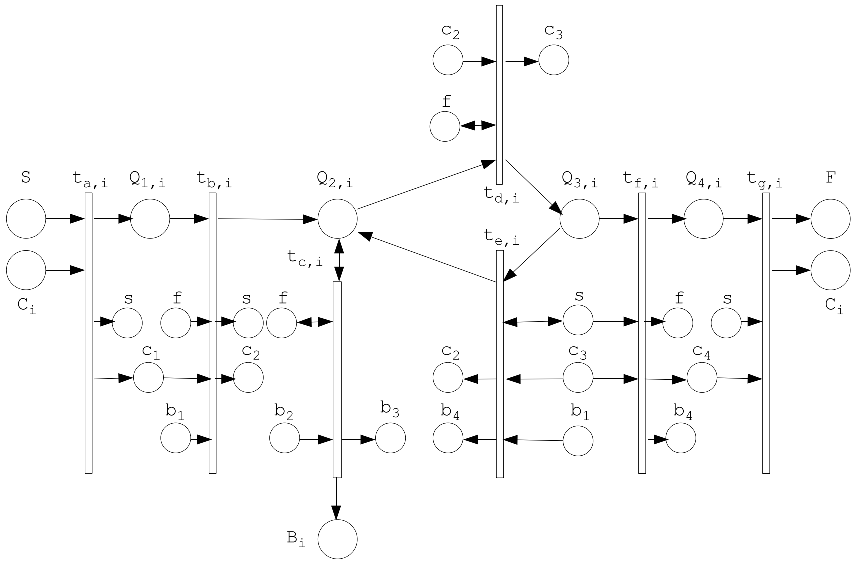

We first exhibit the firing sequence $\sigma_{i,k+1}$ from $S + C_i$ to $F + C_i + e_{k+1}B_i$:

- $S + C_i \xrightarrow{t_{a,i}} Q_{1,i} + s + c_1$
- $Q_{1,i} + s + c_1 \xrightarrow{\sigma_{1,k}} Q_{1,i} + f + c_1 + e_k b_1$ using the inductive hypothesis
- $Q_{1,i} + f + c_1 + e_k b_1 \xrightarrow{t_{b,i}} Q_{2,i} + s + c_2 + (e_k - 1)b_1$
  We now describe a firing sequence from
  $Q_{2,i} + s + c_2 + (e_k - j)b_1 + (j - 1)e_k B_i + (j - 1)b_4$ to
  $Q_{2,i} + s + c_2 + (e_k - j - 1)b_1 + je_k B_i + jb_4$ for $1 \leq j \leq e_k - 1$
  - $Q_{2,i} + s + c_2 + (e_k - j)b_1 + (j - 1)B_i + (j - 1)b_4$
    $\xrightarrow{\sigma_{2,k}} Q_{2,i} + f + c_2 + (e_k - j)b_1 + e_k b_2 + (j - 1)b_4$ using the inductive hypothesis
  - $Q_{2,i} + f + c_2 + (e_k - j)b_1 + e_k b_2 + (j - 1)b_4$
    $\xrightarrow{(t_{c,i})^{e_k}} Q_{2,i} + f + c_2 + (e_k - j)b_1 + e_k b_3 + (j + 1)e_k B_i + (j - 1)b_4$
  - $Q_{2,i} + f + c_2 + (e_k - j)b_1 + e_k b_3 + (j + 1)e_k B_i + (j - 1)b_4$
    $\xrightarrow{t_{d,i}} Q_{3,i} + f + c_3 + (e_k - j)b_1 + e_k b_3 + (j + 1)e_k B_i + (j - 1)b_4$
  - $Q_{3,i} + f + c_3 + (e_k - 1)b_1 + e_k b_3 + (j + 1)e_k B_i + (j - 1)b_4$
    $\xrightarrow{\sigma^{-}_{3,k}} Q_{3,i} + s + c_3 + (e_k - j)b_1 + (j + 1)e_k B_i + (j - 1)b_4$
  - $Q_{3,i} + s + c_3 + (e_k - j)b_1 + (j + 1)e_k B_i + (j - 1)b_4$
    $\xrightarrow{t_{e,i}} Q_{2,i} + s + c_2 + (e_k - j - 1)b_1 + (j + 1)e_k B_i + jb_4$
- After the previous iterations, when reaching $Q_{2,i}+s+c_2+(e_k-1)e_k B_i+(e_k - 1)b_4$, we perform all the steps of the iteration except the last one reaching $Q_{3,i} + s + c_3 + (e_k)^2 B_i + (e_k - 1)b_4 = Q_{3,i} + s + c_3 + e_{k+1}B_i + (e_k - 1)b_4$.

- $Q_{3,i} + s + c_3 + e_{k+1}B_i + (e_k - 1)b_4 \xrightarrow{t_{f,i}} Q_{4,i} + f + c_4 + e_{k+1}B_i + e_k b_4$
- $Q_{4,i} + f + c_4 + e_{k+1}B_i + e_k b_4 \xrightarrow{\sigma^{-}_{4,k}} Q_{4,i} + s + c_4 + e_{k+1}B_i$
- $Q_{4,i} + s + c_4 + e_{k+1}B_i \xrightarrow{t_{g,i}} F + C_i + e_{k+1}B_i$

Let us now prove that any marking reachable from $S_k + C_{i,k} + \alpha_{i_1}B_{i_1,k} + \alpha_{i_2}B_{i_2,k} + \alpha_{i_3}B_{i_3,k}$ ($\{i_1, i_2, i_3\} = \{1, 2, 3, 4\} \setminus \{i\}$) with S_k or F_k marked is either $S_k + C_{i,k} + \alpha_{i_1}B_{i_1,k} + \alpha_{i_2}B_{i_2,k} + \alpha_{i_3}B_{i_3,k}$ or $F_k + C_{i,k} + e_kB_{i,k} + \alpha_{i_1}B_{i_1,k} + \alpha_{i_2}B_{i_2,k} + \alpha_{i_3}B_{i_3,k}$. We first observe on the net above that the tokens contained in a place $B_{j,k}$ are frozen except when place $C_{j,k}$ is marked. Thus w.l.o.g. we assume that $\alpha_1 = \alpha_2 = \alpha_3 = 0$.

So it remains to show that when deviating from the exhibited sequence one cannot reach a marking with S_k marked different from the initial marking or a marking with F_k marked different from the final marking. This is proven by a case study (see [23]). Here we just handle one case since all cases are similar. When reaching marking $Q_{2,i}+f+c_2+(e-j)b_1+e_kb_3+(j+1)e_kB_i+(j-1)b_4$ with $0 \leq e < e_k$, one can fire transition $t_{d,i}$ reaching marking $Q_{3,i} + f + c_3 + (e - j)b_1 + eb_3 + (j + 1)e_kB_i + (j - 1)b_4$. From this marking due to the inductive hypothesis, it is not possible to mark place s in subnet $\mathcal{N}_{k-1}$. Thus transitions $t_{e,i}$ and $t_{f,i}$ are not fireable. So the only possible way to "progress" in $\mathcal{N}_k$ consists to fire the reverse transition $t^{-}_{d,i}$ coming back to the marking $Q_{2,i} + f + c_2 + (e - j)b_1 + e_kb_3 + (j + 1)e_kB_i + (j - 1)b_4$.

The subnet below describes the initial behaviour of the simulating net consisting in filling places $B_{i,n}$ (with $i \in \{1, 2, 3, 4\}$) with e_n tokens and putting a token in q_0 the place corresponding to the initial state of the counter machine.

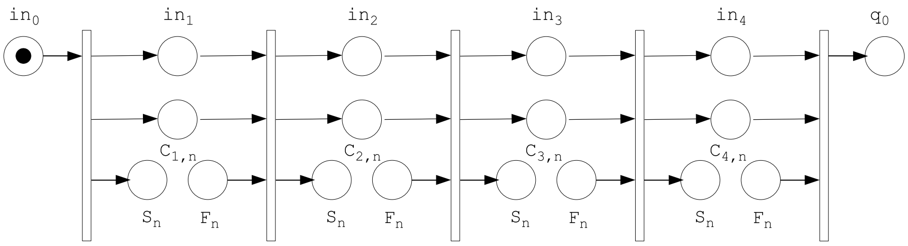

The simulation of an instruction q : if$c_i > 0$ then $c_i{-}{-}$; goto q' else goto q'' is now simply performed by the following subnet. The validity of the zero test is ensured by the assertions about the subnet $\bigcup_{l \leq n} \mathcal{N}_l$. Furthermore it can be proved that reverse transitions of the ones simulating transitions cannot help to mark place q_f where q_f is the final state of the counter machine (see [23] or proposition 12 in [14] for a simple proof of this claim).

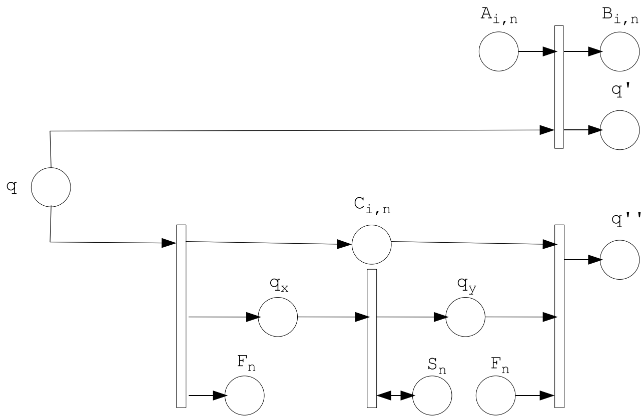

We are now ready to explain the modifications that we bring to the simulating net. For every pair of transitions t and t^-, we add a place p_t input of one of the transitions and output of the other. Thus by construction p_t and $-p_t$ are witnesses for t and t^-. More precisely p_t is the witness of the transition for which it is an output and $-p_t$ is the output of the other transition. Let us examine how these additional places modify the behaviour of the net. Since there is no new transition, firing sequences of the enlarged net are firing sequences of the original one. Thus we only have to care whether the simulation firing sequence is still a firing sequence.

For the transitions not belonging to the subnet $\bigcup_{l\leq n}\mathcal{N}_l$, p_t is an output of t. As the reversed transitions of these transitions do not occur in the simulating sequence, such places cannot disable a transition in the simulating sequence. We now observe that the sequences $\sigma_{i,n}$ and $\sigma_{i,n}^-$ alternate in the simulating sequence, always starting by $\sigma_{i,n}$. Thus in subnet $\mathcal{N}_n$, place p_t is the output of the transition t. Now observe that in sequence $\sigma_{i,n}$ there is an occurrence of sequence $\sigma_{1,n-1}$, e_{n-1} occurrences of $\sigma_{2,n-1}$ followed by $\sigma_{3,n-1}^-$ and then an occurrence of $\sigma_{4,n-1}$. Thus in subnet $\mathcal{N}_{n-1}$, place p_t is the output of the transition t (resp. t^-) when t is $t_{u,i}$ with $u \in \{a,b,c,d,e,f,g\}$ and $i \in \{1,2\}$ (resp. $i \in \{3,4\}$)using notations of the figure. The same pattern of occurrences also happens at lower levels. So more generally, in subnet $\mathcal{N}_k$ with $k < n$, place p_t is the output of the transition t (resp. t^-) when t is $t_{u,i}$ with $u \in \{a,b,c,d,e,f,g\}$ and $i \in \{1,2\}$ (resp. $i \in \{3,4\}$). With this choice, the simulating sequence is still a firing sequence in the enlarged net and the marking to be covered is q_f. □

The complexity of reachability for $\mathsf{\Pi}^2$-nets remains an open issue (indeed the proof of EXPSPACE-hardness does not work for reachability).

5 The subclass of Π^3-nets

In this section, we introduce Π^3-nets, a subclass of product-form Petri nets for which the normalising constant can be efficiently computed. The first subsection defines the subclass; the second one studies its structural properties and the third one is devoted to the computation of the normalising constant.

5.1 Definition and properties

Definition 5.1 (Ordered Π-net)**.** *Consider an integer $n \geq 2$. An n-level* ordered Π-net *is a Π-net $\mathcal{N} = (P, T, W^-, W^+)$ such that:*

1. *$P = \bigsqcup_{1\leq i\leq n} P_i$, $T = \bigsqcup_{1\leq i\leq n} T_i$ and $P_i \neq \emptyset$ for all $1 \leq i \leq n$,*
2. *$\mathcal{M}_i = (P_i, T_i, W^-_{|P_i\times T_i}, W^+_{|P_i\times T_i})$ is a strongly connected state machine,*
3. *$\forall 1 \leq i \leq n$, $\forall t \in T_i$, $\forall p \in P$, ${}^\bullet t(p) > 0$ implies $p \in P_i$ or $p \in P_{i-1}$ ($P_0 = \emptyset$),*
4. *$\forall 2 \leq i \leq n$, $\exists t \in T_i$, $\exists p \in P_{i-1}$ s.t. ${}^\bullet t(p) > 0$,*
5. *$\forall 1 \leq i \leq n$, $\forall t, t' \in T_i$, $({}^\bullet t \cap {}^\bullet t') \cap P_i \neq \emptyset$ implies ${}^\bullet t = {}^\bullet t'$.*

We call $\mathcal{M}_i$ the level i state machine. The elements of P_i (resp. T_i) are level i places (resp. transitions). The complexes ${}^\bullet t$ with $t \in T_i$ are level i complexes.

By weak reversibility, the constraints 3, 4, and 5 also apply to the output bags $t^\bullet$. An *ordered* Π*-net* is a sequence of strongly connected state machines. Connections can only be made between a level i transition and a level $(i-1)$ place (points 1, 2, 3). By construction, an *ordered* Π*-net* is connected (point 4). For $i > 1$, each level i place belongs to one and only one level i complex (point 5). An example of ordered Π-net can be found on figure 7.

Lemma 5.2. *The reaction net of $\mathcal{N}$ is isomorphic to the disjoint union of state machines $\mathcal{M}_i$. Consequently, a T-semi-flow of $\mathcal{M}_i$ is also a T-semi-flow of $\mathcal{N}$. If a transition of T_i is enabled by a reachable marking then every transition of T_i is live.*

Proof. Consider the mapping f which maps each complex $t^\bullet$, $t \in T_i$, to p the output place of t in P_i. By construction of ordered Π-nets, f is a bijection from $\mathcal{C}$ to P. Moreover, each arc $c_1 \to c_2$ of the reaction graph corresponds to the transition $t = f(c_1)^\bullet = {}^\bullet f(c_2)$. This proves the first point of the lemma.
To prove the second point, recall that for a state machine, the T-semi-flows correspond to circuits of the Petri net graph. From this and from the first point, a T-semi-flow of $\mathcal{M}_i$ defines a circuit of the reaction graph of $\mathcal{N}$, which yields a T-semi-flow of $\mathcal{N}$.
The set of transitions T_i is the set of transitions occurring in a component of the reaction graph. The third point follows. □

An ordered Π-net may be interpreted as a multi-level system. The transitions represent jobs or events while the tokens in the places represent resources or constraints. A level i job requires resources from level $(i-1)$ and relocates these resources upon completion. On the contrary, events occurring in level $(i-1)$ may make some resources unavailable, hence interrupting activities in level i. The dependency of an activity on the next level is measured by *potentials*, defined as follows.

Definition 5.3 (Interface, potential)**.** *A place $p \in P_i$, $1 \le i \le n-1$, is an* interface place *if $p \in t^\bullet$ for some $t \in T_{i+1}$. For a place $p \in P_i$, $2 \le i \le n$, and a place $q \in P_{i-1}$, set:*

$$pot(p,q) = \begin{cases} t^\bullet(q) & \text{if } p \text{ and } q \text{ have a common input transition } t \in T_i \\ 0 & \text{otherwise.} \end{cases}$$

The potential *of a place $p \in P_i$, $2 \le i$, is defined by:*

$$pot(p) = \sum_{q \in P_{i-1}} pot(p,q)\,.$$

By convention, $pot(p) = 0$ for all $p \in P_1$.

By the definition of ordered Π-nets, the quantity $t^\bullet(q)$ does not depend on the choice of t, so the potential is well-defined. Indeed, by weak reversibility, the constraint 5 also applies to the output bags $t^\bullet$.

Example. The Petri net in Figure 7 is a 3-level ordered Π-net. The potentials are written in parentheses. To keep the figure readable, the arcs between the place p_1 and the level 2 transitions are not shown.

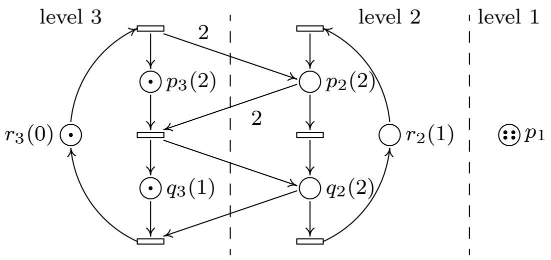

Figure 7: Ordered Π-net.

Definition 5.4 (Marking witness)**.** *The* marking witness *of a marking m, denoted by $\widetilde{m}$, is defined as follows. For all $i \le n$ and $p \in P_i$,*

$$\widetilde{m}(p) = m(p) + \sum_{j=1}^{n-i} \Big((-1)^j \sum_{\substack{r_1 \in P_{i+1} \\ r_j \in \dot{P}_{i+j}}} m(r_j) \big(\prod_{k=1}^{j-1} pot(r_{k+1}, r_k) \big) pot(r_1, p) \Big)\,. \qquad (5.1)$$

Remark. Note that a marking witness is not necessarily non-negative. It can be showed by induction that:
$\forall p \in P_n\,, \widetilde{m}(p) = m(p)$ and $\forall p \in P_i\,, i < n\,, \widetilde{m}(p) = m(p) - \sum_{r \in P_{i+1}} \widetilde{m}(r) pot(r,p)$

Lemma 5.5. *Let m, m' be two vectors such that $m' = m + W(t)$ for some $t \in T_i$ $(1 \leq i \leq n)$. Let p_1 and p_2 denote the input place and the output place of t in P_i, respectively. Then for every place p:*

$$\widetilde{m}'(p) = \widetilde{m}(p) - 1 \textit{ if } p \textit{ is } p_1, \quad \widetilde{m}(p) + 1 \textit{ if } p \textit{ is } p_2, \quad \widetilde{m}(p) \textit{ otherwise.} \tag{5.2}$$

Proof. Since m and m' have the same restriction on $\cup_{j>i} P_j$, we have $\widetilde{m}'(p) = \widetilde{m}(p)\ \forall p \in (\cup_{j \geq i} P_j) \setminus \{p_1, p_2\}$. It follows that $\widetilde{m}'(p_1) - \widetilde{m}(p_1) = m'(p_1) - m(p_1) = -1$ and $\widetilde{m}'(p_2) - \widetilde{m}(p_2) = m'(p_2) - m(p_2) = 1$.
For $p \in P_{i-1} \cap t^\bullet$, we have $m'(p) - m(p) = pot(p_2, p) - pot(p_1, p)$, hence

$$\begin{aligned}\widetilde{m}'(p) - \widetilde{m}(p) &= m'(p) - m(p) - \big[(\widetilde{m}'(p_1) - \widetilde{m}(p_1)) pot(p_1, p) \\ &\quad + (\widetilde{m}'(p_2) - \widetilde{m}(p_2)) pot(p_2, p)\big] \\ &= 0\end{aligned}$$

Similarly, $\widetilde{m}'(p) - \widetilde{m}(p) = 0$ for $p \in P_{i-1} \cap {}^\bullet t$.
For all other places, $m'(p) = m(p)$ and $\widetilde{m}'(r) = \widetilde{m}(r)\ \forall r$ s.t. $pot(r,p) \neq 0$, thus $\widetilde{m}'(p) = \widetilde{m}(p)$. □

The above lemma applies in particular when m and m' are markings such that $m \xrightarrow{t} m'$. Equations (5.2) look like the equations for witnesses. Since each level i complex contains exactly one level i place, one guesses that every complex admits a witness, i.e. that $\mathcal{N}$ is a Π^2-net. This is confirmed by the next proposition.

Proposition 5.6. *Let B denote the $P \times P$ integer matrix of the linear transformation $m \mapsto \widetilde{m}$ defined by (5.1). For $p \in P_i$, the line vector $B(p)$ is a witness for the i-level complex containing p. In particular, $\mathcal{N}$ is a Π^2-net.*

Proof. Denote by $A \in \mathbb{Z}(\mathcal{C} \times T)$ the incidence matrix of the reaction graph. From Lemma 5.5, we have:

$$m \xrightarrow{t} m' \implies \widetilde{m}' - \widetilde{m} = A(t)\,.$$

We have to show that $BW(t) = A(t)\ \forall t \in T$. Indeed, let m and m' be two markings such that $m \xrightarrow{t} m'$, we have: $BW(t) = B(m' - m) = \widetilde{m}' - \widetilde{m} = A(t)$. □

Lemma 5.5 allows to derive relevant S-semi-flows of $\mathcal{N}$ and S-invariants.

Corollary 5.7. *Let m_0 be the initial marking of $\mathcal{N}$. We have:*
$$\forall m \in \mathcal{R}(m_0), \quad \forall i \in \{1, \ldots, n\}, \qquad \widetilde{m}(P_i) = \widetilde{m}_0(P_i)$$
More generally, for all i, the vector $v_i = \sum_{p \in P_i} B(p)$ is a S-semi-flow of $\mathcal{N}$.

Using this corollary, it can be shown that an ordered Π-net is bounded.

Example. Consider the ordered Π-net in Figure 7 with the initial marking $m_0 = p_3 + q_3 + r_3 + 4p_1$. The marking witness of m_0 is $\widetilde{m_0} = p_3 + q_3 + r_3 - 2p_2 - q_2 + 10p_1$. Any reachable marking m satisfies the invariants:

$$m(P_3) = 3$$
$$m(P_2) - 2m(p_3) - m(q_3) = -3$$
$$m(p_1) - 2m(p_2) - 2m(q_2) - m(r_2) + 4m(p_3) + 2m(q_3) = 10$$

We shown that $\{v_i, 1 \leq i \leq n\}$ is a basis of the S-semi-flows of $\mathcal{N}$.

Proposition 5.8. *Let v be an S-semi-flow of $\mathcal{N}$, i.e. $v.W = 0$. There exist unique rational numbers $a_1, \ldots, a_n$ such that $v = \sum_{i=1}^{n} a_i v_i$.*

Proof. The matrix B is a $P \times P$ unit lower triangular matrix, so it is invertible.
We have:

$$v.W = 0 \implies (v.B^{-1})(BW) = 0 \implies (v.B^{-1})A = 0\,,$$

hence $v.B^{-1}$ is an S-semi-flow of the disjoint union of the state machines $\mathcal{M}_i$. But since a state machine's only S-semi-flows are $a(1, \ldots, 1)$, $a \in \mathbb{Q}$, there exist rational numbers $a_1, \ldots, a_k$ such that

$$v.B^{-1} = \sum_{i=1}^{n} a_i w_i\,, \tag{5.3}$$

where $w_i \in \mathbb{Q}^P$ are defined by $w_i(p) = \mathbb{1}_{P_i}(p)$.
Right-multiplying both sides of (5.3) by B, we get $v = \sum_{i=1}^{n} a_i v_i$.
The independence of the set $\{v_i\,, 1 \leq i \leq n\}$ follows from the fact that the vectors $v_i B^{-1}$ have non-empty disjoint supports. □

We now consider only ordered Π-nets in which the interface places in P_i have maximal potential among the places of P_i. From the technical point of view, this assumption is crucial for the reachability set analysis presented later. From the modelling point of view, it is a reasonable restriction. Consider the multi-level model, the assumption means that during the executions of level i jobs, the level $(i-1)$ is idle, therefore the amount of available resource is maximal.

Definition 5.9 (Π^3-net)**.** *An ordered Π-net $\mathcal{N}$ is a Π^3-net if:*

$$\forall i, \forall p \in P_i : p \in {}^{\bullet}T_{i+1} \implies pot(p) = max\{pot(q), q \in P_i\}\,.$$

5.2 The reachability set

From now on, $\mathcal{N}$ is a n-level Π^3-net with $\mathcal{M}_1, \ldots, \mathcal{M}_n$ being its state machines.

Definition 5.10 (Minimal marked potential)**.** *Consider* $i \in \{2, \ldots, n\}$. *The* level i minimal potential marked by m *is:*

$$\varphi_i(m) = \begin{cases} max\{pot(p), p \in P_i\} & \text{if } m(P_i) = 0\,, \\ min\{pot(p), p \in P_i, m(p) > 0\} & \text{if } m(P_i) > 0\,. \end{cases}$$

The next lemma gives a necessary condition for reachability.

Lemma 5.11. *If* $\varphi_i(m) \le m(P_{i-1})$ *then* $\varphi_i(m') \le m'(P_{i-1})$ *for all* $m' \in \mathcal{R}(m)$.

Proof. W.l.o.g., assume that $m \xrightarrow{t} m'$.
First, suppose that $t \notin T_i$. If $t \notin T_{i+1}$ then firing t does not modify the marking on P_i, so $\varphi_i(m') = \varphi_i(m)$. If $t \in T_{i+1}$, firing t either leaves the marking of P_i unchanged or moves tokens between places of maximal potential in P_i; in both cases $\varphi_i(m') = \varphi_i(m)$. Since $t \notin T_i$, $m'(P_{i-1}) = m(P_{i-1})$. So $\varphi_i(m') \le m'(P_{i-1})$ if $t \notin T_i$.
Now consider $t \in T_i$, let p and q be the input and output places of t in P_i. We have $\varphi_i(m') \le pot(q) \le m(P_{i-1}) - pot(p) + pot(q) = m'(P_{i-1})$. □

We now define the partial liveness and partial reachability.

Definition 5.12 (i-reachability set, i-liveness)**.** *Let* m *be a marking. The* i*-reachability set of* m*, denoted by* $\mathcal{R}_i(m)$*, is the set of all markings reachable from* m *by a firing sequence consisting of transitions in* $\bigcup_{1 \le j \le i} T_j$. *We say that* m *is* i-live *if for any transitions* t *in* $\bigcup_{1 \le j \le i} T_j$*, there exists a marking in* $\mathcal{R}_i(m)$ *which enables* t. *By convention,* $\mathcal{R}_0(m) = \{m\}$ *and every marking is* 0*-live.*

The i-live markings are characterised by the following proposition.

Proposition 5.13. *A marking* m *is* i*-live if and only if it satisfies the following inequalities, called the* i-condition*:*

$$m(P_i) > 0 \wedge \forall 2 \le j \le i \,:\, m(P_{j-1}) \ge \varphi_j(m) \tag{5.4}$$

If m *satisfies the* i*-condition then for every* $p, q \in P_i$ *such that* $p \neq q$, $m(p) > 0$ *and* $pot(p) \le m(P_{i-1})$*, there exists* $m' \in \mathcal{R}_i(m)$ *such that:*

$$m'(p) = m(p) - 1\,,\; m'(q) = m(q) + 1\,,\; \forall r \in P_i \setminus \{p, q\},\; m'(r) = m(r)\,. \tag{5.5}$$

A marking is live if and only if it satisfies the n*-condition.*

Proof. Consider an i-live marking m. For any $j \le i$, there is a marking $m' \in \mathcal{R}_i(m)$ which enables a transition of T_j. This marking satisfies $\varphi_j(m') \le m'(P_{j-1})$. By (weak) reversibility, $m \in \mathcal{R}(m')$, so $\varphi_j(m) \le m(P_{j-1})$ (Lemma 5.11). Since the number of tokens in P_i is the same for all the markings of $\mathcal{R}_i(m)$, $m(P_i) > 0$ (otherwise, the transitions of T_i would be dead).

We prove the reverse direction and the second part of the proposition by induction on $i \geq 1$, i.e. :
If m satisfies the i-condition then:
(1) for every $p, q \in P_i$ such that $p \neq q$, $m(p) > 0$ and $pot(p) \leq m(P_{i-1})$, there exists $m' \in \mathcal{R}_i(m)$ such that: $m'(p) = m(p) - 1\,,\ m'(q) = m(q) + 1\,,\ \forall r \in P_i \setminus \{p, q\},\ m'(r) = m(r)\,.$
(2) m is i-live.

The case $i = 1$ is trivial.
Suppose that the claim has been proven for all $j \leq i - 1$. Let m be a marking which satisfies the i-condition. Consider two cases: $pot(p) = 0$ and $pot(p) > 0$.
If $pot(p) = 0$ then the output transitions of p are enabled by m. For any arbitrary $q \neq p$, fire the transitions along a path from p to q in T_i, we obtain a marking m' satisfying (5.5). So we have proved assertion (1). Now choose some q such that $pot(q) > 0$ (there is at least one). Then $m'(P_{i-1}) \geq pot(q) > 0$. By the induction hypothesis, m' is $(i-1)$-live. Moreover, m' enables the output transitions of q. Hence m' is i-live, which implies m is i-live.
If $pot(p) > 0$ then $m(P_{i-1}) > 0$, hence m is $(i-1)$-live by the induction hypothesis. It remains to find a marking in $\mathcal{R}_{i-1}(m)$ which enables the output transitions of p. If for all $r \in P_{i-1}$, $m(r) \geq pot(p, r)$ then choose m. Otherwise, choose a marked place q of P_{i-1} such that $pot(q) \leq m(P_{i-2})$ and a level $(i-1)$ interface place q', then apply the induction hypothesis on (5.5) to find $m_1 \in \mathcal{R}_{i-1}(m)$ such that $m_1(q) = m(q) - 1$, $m_1(q') = m(q') + 1$ and $m_1(r) = m(r)$ for every other places r of P_{i-1}. We have $\varphi_{i-1}(m_1) = \max\{pot(r), r \in P_{i-1}\}$. Now starting from m_1, repeat the following procedure:

- Step 1: Find two place r_1, r_2 in P_{i-1} such that $\bar{m}(r_1) < pot(p, r_1)$ and $\bar{m}(r_2) > pot(p, r_2)$, $\bar{m}$ denoting the current marking.
- Step 2: Use the induction hypothesis on (5.5) to find $\bar{m}' \in \mathcal{R}_{i-1}(\bar{m})$ such that $\bar{m}'(r_1) = \bar{m}(r_1) + 1$, $\bar{m}'(r_2) = \bar{m}(r_2) - 1$ and $\bar{m}'(r) = \bar{m}(r)$ for all $r \in P_{i-1} \setminus \{r_1, r_2\}$.

All the intermediate markings are $(i-1)$-live. Since $\bar{m}(P_{i-1}) \geq pot(p)$, if there exists $r_1 \in P_{i-2}$ such that $\bar{m}(r_1) < pot(p, r_1)$ then there exists $r_2 \in P_{i-2}$ such that $\bar{m}(r_2) > pot(p, r_2) \geq 0$ as well. Because the interface places have maximal potential, at the beginning of each iteration, we always have $\varphi_{i-1}(\bar{m}) = \max\{pot(r), r \in P_{i-1}\}$, hence $pot(r_2) \leq \varphi_{i-1}(\bar{m}) \leq \bar{m}(P_{i-2})$. Each iteration strictly diminishes the number of "missing" tokens in the places of P_{i-1} synchronised with p, so the procedure eventually stops at a marking m_2 such that $m_2(r) \geq pot(p, r)$ for every place $r \in P_{i-1}$. This marking enables the output transitions of p. □

Example: The ordered Π-net in Figure 7 is a Π^3-net. Consider two markings: $m_1 = p_3 + q_3 + r_3 + 4p_1$ and $m_2 = 3q_3 + 4p_1$. These markings agree on all the S-invariants, but only m_1 satisfies the 3-condition. It is easy to check that m_1 is live while m_2 is dead.

We conclude this subsection by showing that the reachability problem for Π^3-nets can be efficiently decided as well.

Theorem 5.14. *Suppose that the initial marking m_0 is live. Then the reachability set $\mathcal{R}(m_0)$ coincides with the set $\mathcal{S}(m_0)$ of markings which satisfy the n-condition and agree with m_0 on the S-invariants given by Corollary 5.7.*

Proof. The inclusion $\mathcal{R}(m_0) \subset \mathcal{S}(m_0)$ is the combination of the results of Corollary 5.7 and Proposition 5.13.
To prove the converse, we look for a marking which is reachable from every marking of $\mathcal{S}(m_0)$. Let p_j, $1 \leq j \leq n$, be a place of maximal potential of P_j, that is, $pot(p_j) = \max\{pot(p), p \in P_j\}$. Let m'_0 denote the unique marking in $\mathcal{S}(m_0)$ such that $m'_0(p) = 0$ for every $p \notin \{p_1, \ldots, p_n\}$. Consider an arbitrary marking m in $\mathcal{S}(m_0)$. We prove by a reverse induction on $i \leq n$ and by using the second part of Proposition 5.13 that there exists a marking $m' \in \mathcal{R}(m)$ such that $m'(p) = 0 \ \forall p \notin \{p_1, \ldots, p_n\}$. The inductive claim is:

$$\begin{array}{c}\text{There exists a marking } m'_i \in \mathcal{R}(m) \text{ such that} \\ \forall p \in \cup_{i \leq j \leq n} P_j \setminus \{p_i, \ldots, p_n\} \ m'_i(p) = 0 \\ \text{and } m'_i \text{ satisfies the } i-1 \text{ condition.}\end{array}$$

Let us address the basis case $i = n$. Assume that there exists $p \neq p_n$ such that $m_0(p) > 0$. Using proposition 5.13, we move a token from p to p_n. Furthermore by lemma 5.11, the n-condition is still satisfied. Iterating this process, we obtain a marking m'_n such that $\forall p \in P_n \setminus \{p_n\} \ m'_n(p) = 0$ and the n-condition is still satisfied. The inductive case is similar by observing that the sequence that moves the tokens of P_i does not use transitions of T_j for $j > i$.
Since m' is also an element of $\mathcal{S}(m_0)$, $m' = m'_0$. So m'_0 is reachable from every marking in $\mathcal{S}(m_0)$. By (weak) reversibility, every marking in $\mathcal{S}(m_0)$ is reachable from m'_0. So $\mathcal{S}(m_0) \subset \mathcal{R}(m'_0) = \mathcal{R}(m_0)$. □

5.3 Computing the normalising constant

The normalising constant of a product-form Petri net (see Section 2.1) is $G = \sum_m \mathbb{1}_{m \in \mathcal{R}(m_0)} \prod_{p \in P} u_p^{m(p)}$. It is in general a difficult task to compute G, as can be guessed from the complexity of the reachability problem. However, efficient algorithms may exist for nets with a well-structured reachability set. Such algorithms were known for Jackson networks [25] and the *S-invariant reachable* Petri nets defined in [11]. We show that is is also the case for the class of live Π^3-nets which is strictly larger than the class of Jackson networks (which correspond to 1-level ordered nets) and is not included in the class of S-invariant reachable Petri nets.

Suppose that m_0 is a live marking. Suppose that the places of each level are ordered by increasing potential: $P_i = \{p_{i1}, \ldots, p_{ik_i}\}$ such that $\forall 1 \le j < k_i$, $pot(p_{ij}) \le pot(p_{i(j+1)})$.

Let V denote the $n \times P$-matrix the i-th row of which is the S-invariant v_i defined in Corollary 5.7. For $1 \le i \le n$, set $C_i = v_i m_0 = \widetilde{m}_0(P_i)$. Then the reachability set consists of all n-live markings m such that $Vm = {}^t(C_1, \ldots, C_n)$.

For $1 \le i \le n$, $1 \le j \le k_i$ and $c_1, \ldots, c_i \in \mathbb{Z}$, define $E(i, j, c_1, \ldots, c_i)$ as the set of markings m such that

$$\begin{cases} m(p_{i\nu}) = 0 \text{ for all } \nu > j \\ Vm = {}^t(c_1, \ldots, c_i, 0 \ldots, 0) \\ \varphi_\nu(m) \le m(P_{\nu-1}) \text{ for all } 2 \le \nu \le i\,. \end{cases}$$

The elements of $E(i, j, c_1, \ldots, c_i)$ are the markings which satisfy the second part of the i-condition and the S-invariants constraints $(c_1, \ldots, c_i, 0, \ldots, 0)$ and concentrate tokens in $P_1, \ldots, P_{i-1}$ and $\{p_{i1}, \ldots, p_{ij}\}$.

With each $E(i, j, c_1, \ldots, c_i)$ associate

$$G(i, j, c_1, \ldots, c_i) = \pi(E(i, j, c_1, \ldots, c_i)) = \sum \prod_{p \in P} u_p^{m(p)}$$

the sum being taken over all $m \in E(i, j, c_1, \ldots, c_i)$.

We propose to compute $G(n, k_n, C_1, \ldots, C_n)$ by dynamic programming. It consists in breaking each $G(i, j, c_1, \ldots, c_i)$ into smaller sums. This corresponds to a partition of the elements of $E(i, j, c_1, \ldots, c_i)$ by the number of tokens in p_{ij}.

Proposition 5.15. *Let be given $E = E(i, j, c_1, \ldots, c_i)$. If $c_i < 0$ then $E = \emptyset$. If $c_i \ge 0$ then for every non-negative integer a:*

1. *If $a > c_i$ then $E \cap \{m | m(p_{ij}) = a\} = \emptyset$.*
2. *If $a < c_i$ and $j = 1$ then $E \cap \{m | m(p_{ij}) = a\} = \emptyset$.*
3. *If $a < c_i$ and $j \ge 2$ then $E \cap \{m | m(p_{ij}) = a\} = \{m + ap_{ij} \mid m \in E(i, j-1, c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(a\ p_{ij}))\}$.*
4. *If $a = c_i$ and $i = 1$ then $E \cap \{m | m(p_{ij}) = a\} = \{c_1 p_{1j}\}$.*
5. *If $a = c_i$ and $i > 1$ then $E \cap \{m | m(p_{ij}) = a\} = \{\ m + ap_{ij} \mid m \in E(i-1, k_{i-1}, c_1 - v_1(ap_{ij}), \ldots, c_{i-1} - v_{i-1}(ap_{ij}))\}$.*

Proof. Suppose that $E \ne \emptyset$. Let m be an element of E such that $m(p_{ij}) = a$. We have $m(P_i) = c_i$, so $a \le c_i$. Moreover, if $m(p_{ij}) < m(P_i)$ then m must mark some place $p_{i\nu}$ with $\nu < j$, so $j \ge 2$. These prove the first and the second cases.

The fourth case is trivial.

Let us address the third case, we have to show that:

$$\forall m \in E \text{ s.t. } m(p_{ij}) = a, (m - ap_{ij}) \in E(i, j-1, c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(ap_{ij})) \tag{5.6}$$

$$\forall m' \in E(i, j-1, c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(ap_{ij})), (m' + ap_{ij}) \in E \tag{5.7}$$

The values $c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(ap_{ij})$ are obtained by:

$$\begin{aligned} Vm &= {}^t(c_1, \ldots, c_i, 0, \ldots, 0) \\ \Longleftrightarrow \; V(m - ap_{ij}) &= {}^t(c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(ap_{ij}), 0, \ldots, 0)\,. \end{aligned}$$

We have to show that $\varphi_\nu(m - ap_{ij}) \le (m - ap_{ij})(P_{\nu-1})$ $\forall 2 \le \nu \le i$ and $\varphi_\nu(m' + ap_{ij}) \le (m' + ap_{ij})(P_{\nu-1})$ $\forall 2 \le \nu \le i$.
Since m and $(m - ap_{ij})$ only differ at p_{ij}, it suffices to show that $\varphi_i(m - ap_{ij}) \le (m - ap_{ij})(P_{i-1})$. Indeed, $\varphi_i(m - ap_{ij}) = \varphi_i(m)$ because both markings mark some $p_{i\nu}$ with $\nu < j$, and $(m - ap_{ij})(P_{i-1}) = m(P_{i-1})$ because the two markings are identical on P_{i-1}.
Similarly, given $m' \in E(i, j-1, c_1 - v_1(ap_{ij}), \ldots, c_i - v_i(ap_{ij}))$, to prove (5.7), it suffices to show that $\varphi_i(m' + ap_{ij}) \le (m' + ap_{ij})(P_{i-1})$. Indeed, $(m' + ap_{ij})(P_{i-1}) = m'(P_{i-1}) \le \varphi_i(m') \le \varphi_i(m' + ap_{ij})$.

The fifth case is similar. It suffices to show that $\varphi_{i-1}(m - ap_{ij}) \le (m - ap_{ij})(P_{i-2})$ and $\varphi_i(m' + ap_{ij}) \le (m' + ap_{ij})(P_{i-1})$. The first inequality is immediate since $(m - ap_{ij})$ is the restriction of m on $\bigcup_{1 \le \nu \le i-1} P_\nu$. To prove the second one, note that $(m' + ap_{ij})(P_{i-1}) = m'(P_{i-1}) = c_{i-1} - v_{i-1}(ap_{ij}) = m(P_{i-1})$ and $\varphi_i(m' + ap_{ij}) = \varphi_i(m)$. □

The proposition 5.15 induces the following relations between the sums $G(i, j, c_1, \ldots, c_i)$.

Corollary 5.16. *If $c_i < 0$ then $G(i, j, c_1, \ldots, c_i) = 0$. If $c_i \ge 0$ then:*

- *Case $2 \le i \le n$, $2 \le j \le k_i$:*

$$\begin{aligned} G(i, j, c_1, \ldots, c_i) = &\sum_{\nu=0}^{c_i - 1} u_{p_{ij}}^{\nu} G(i, j-1, c_1 - v_1(\nu p_{ij}), \ldots, c_i - v_i(\nu p_{ij})) \\ &+ u_{p_{ij}}^{c_i} G(i-1, k_{i-1}, c_1 - v_1(c_i p_{ij}), \ldots, c_{i-1} - v_{i-1}(c_i p_{ij}))\,. \end{aligned}$$

- *Case $2 \le i \le n$, $j = 1$:*

$$G(i, 1, c_1, \ldots, c_i) = u_{p_{i1}}^{c_i} G(i-1, k_{i-1}, c_1 - v_1(c_i p_{i1}), \ldots, c_{i-1} - v_{i-1}(c_i p_{i1}))\,.$$

- *Case $i = 1$, $j \ge 2$:* $G(1, j, c_1) = \sum_{\nu=0}^{c_1 - 1} u_{p_{1j}}^{\nu} G(1, j-1, c_1 - \nu) + u_{p_{1j}}^{c_1}$.
- *Case $i = 1$, $j = 1$:* $G(1, 1, c_1) = u_{p_{11}}^{c_1}$.

Complexity. Since $i \leq n$, $j \leq K = \max\{k_1, \ldots, k_n\}$, the number of evaluations is bounded by $n \times K \times \gamma$, where γ upper bounds the c_i's. Let α denote the global maximal potential. From (5.1), we obtain $\gamma = \mathcal{O}(m_0(P)K^n\alpha^n)$. So the complexity of a dynamic programming algorithm using Cor. 5.16 is $\mathcal{O}(m_0(P)nK^{n+1}\alpha^n)$, i.e. pseudo-polynomial for a fixed number of state machines.

6 Perspectives

This work has several perspectives. First, we are interested in extending and applying our rules for a modular modelling of complex product-form Petri nets. We also want to obtain characterisation of product-form Petri nets when stochastic Petri nets are equipped with infinite-server policy. Then we want to validate the formalism of Π^3-nets showing that it allows to express standard patterns of distributed systems. We plan to implement analysis of Π^3-nets and integrate it into a tool for stochastic Petri nets like Great-SPN [8]. Finally we conjecture that reachability is EXPSPACE-complete for Π^2-nets and we want to establish it.

Acknowledgements. We would like to thank the anonymous referees whose numerous and pertinent suggestions have been helpful in preparing the final version of the paper.

References

[1] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets.* John Wiley & Sons, 1995.

[2] F. Baccelli, G. Cohen, G.J. Olsder, and J.P. Quadrat. *Synchronization and Linearity.* John Wiley & Sons, New York, 1992.

[3] S. Balsamo, P. G. Harrison and A. Marin. Methodological construction of product-form stochastic Petri nets for performance evaluation *Journal of Systems and Software*, available on line, to appear.

[4] S. Balsamo and A. Marin. Performance engineering with product-form models: efficient solutions and applications *Proceedings of the second joint WOSP/SIPEW international conference on Performance engineering (ICPE'11)*, ACM publisher, pp. 437-448, Karlsruhe, Germany, 2011.

[5] S. Balsamo, A. Marin. Composition of product-form Generalized Stochastic Petri Nets: a modular approach. *Proc. ESM 2009, Eurosis*

23rd European Simulation and Modelling Conference, United Kingdom, Leicester, October 2009.

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

[7] R. J. Boucherie, M. Sereno. On closed support T-invariants and traffic equations. *Journal of Applied Probability*, (35): 473–481, 1998.

[8] G. Chiola, G. Franceschinis, R. Gaeta, M. Ribaudo. GreatSPN 1.7: Graphical Editor and Analyzer for Timed and Stochastic Petri Nets. *Performance Evaluation* 24(1-2): 47-68 (1995)

[9] J. Desel and J. Esparza *Free Choice Petri Nets*, volume 40 of *Cambridge Tracts Theoret. Comput. Sci.* Cambridge Univ. Press, 1995.

[10] J. Esparza and M. Nielsen. Decidability issues for Petri nets - a survey. *Journal of Informatik Processing and Cybernetics*, 30(3):143-160, 1994.

[11] J.L. Coleman, W. Henderson, P.G. Taylor. Product form equilibrium distributions and a convolution algorithm for stochastic Petri nets. *Performance Evaluation*, 26(3):159–180, September 1996.

[12] J. Esparza. Reduction and Synthesis of Live and Bounded Free Choice Petri Nets. *Information and Computation*, 114(1):50–87, 1994

[13] M. Feinberg. Lectures on chemical reaction networks. *At the Math. Research Center, Univ. Wisconsin, 1979.* Available online at: `http://www.che.eng.ohio-state.edu/~feinberg/LecturesOnReactionNetworks`.

[14] S. Haddad, P. Moreaux, M. Sereno, M. Silva Product-form and stochastic Petri nets: a structural approach. *Performance Evaluation*, 59: 313-336, 2005.

[15] S. Haddad, J. Mairesse, H.-T. Nguyen Synthesis and analysis of product-form Petri nets. *Petri Nets 2011*, LNCS 6709, 288-307, 2011.

[16] P. G. Harrison and L. M. Catalina Hierarchically constructed Petri-nets and product-forms. *Proceedings of the 5th International ICST Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS'11)*, ICST, pp. 101-110, Paris, France, 2011.

[17] W. Henderson, D. Lucic, P.G. Taylor. A net level performance analysis of stochastic Petri nets. *Journal of Australian Mathematical Soc. Ser. B*, 31:176–187, 1989.

[18] J. R. Jackson. Jobshop-like Queueing Systems. *Management Science*, 10(1): 131–142, 1963.

[19] F. Kelly. *Reversibility and Stochastic Networks*. Wiley, New-York, 1979.

[20] A. A. Lazar, T. G.Robertazzi. Markovian Petri Net Protocols with Product Form Solution. *Proc. of INFOCOM 87*, pp. 1054–1062, San Francisco, CA, USA, 1987.

[21] M. Li, N. D. Georganas. Parametric Analysis of Stochastic Petri Nets, *Fifth International Conference on Modelling and Tools for Computer Performance Evaluation*, Torino, Italy, 1991,

[22] J. Mairesse, H-T. Nguyen. Deficiency Zero Petri Nets and Product Form. *Petri Nets 2009*, LNCS 5606, 103-122, 2009.

[23] E. Mayr, A. Meyer. The complexity of the word problem for commutative semigroups an polynomial ideals. *Advances in Math*, 46 (1982), 305-329.

[24] C. Papadimitriou. Computational Complexity. *Addison Wesley*, 1994.

[25] M. Reiser, S.S. Lavenberg. Mean Value Analysis of Closed Multichain Queueing Networks. *Journal of the ACM*, 27(2): 313-322, 1980.