Properties and Steady-State Performance Bounds for Petri Nets with Unique Repetitive Firing Count Vector

J. Campos*
Dpto. de Ingeniería Eléctrica e Informática
Universidad de Zaragoza, Spain.

G. Chiola†
Dipartimento di Informatica
Università di Torino, Italy.

M. Silva*
Dpto. de Ingeniería Eléctrica e Informática
Universidad de Zaragoza, Spain.

Abstract

The problem of computing both upper and lower bounds for the steady-state performance of timed and stochastic Petri nets is studied. In particular, Linear Programming problems defined on the incidence matrix of underlying Petri net are used to compute bounds for the throughput of transitions for live and bounded nets with a unique possibility of steady-state behaviour. These classes of nets are defined and their characteristics are studied. The bounds proposed here depend on the initial marking and the mean values of the delays but not on the probability distributions (thus including both the deterministic and the stochastic cases); moreover they can be computed also for non-ergodic models. Connections between results and techniques typical of qualitative and quantitative analysis of Petri models are stressed.

1 Introduction

In this paper, which is a continuation of the companion paper [CCCS89], we study the possibility of obtaining (upper and lower) *bounds* on the steady-state performance of a Petri net model, and we restrict the analysis to some net classes that are characterized by having only one possible steady-state behaviour. In particular we study the throughput of transitions, defined as the average number of firings per unit time. From this quantity, applying Little's formula it is possible to derive all average performance estimates of the model.

In a Petri net there is an obvious relation between the concepts of steady-state behaviour and that of repeatable firing sequences: sequences of transitions that are repeatable only a finite number of times cannot contribute to the steady-state performance of the model. In the particular case of safe Marked Graphs, this relation can be made evident by looking at the occurrence nets that describe their behaviours [WT85].

Here we consider live bounded connected nets (thus strongly connected) which are either decision-free or such that the decision policy at effective conflicts is not relevant from the performances point of view (i.e. the decisions do not affect the steady-state performance). A characteristics of these nets is the existence of a unique firing count vector $\vec{\sigma}_R$ associated with all *repetitive sequences* with firing count vectors of the form $\vec{\sigma}$ (such that $C \cdot \vec{\sigma} \geq 0$, in order to be able to repeat the sequence for an arbitrary number of times), i.e. $\vec{\sigma} = k \cdot \vec{\sigma}_R$ with $k \in \mathbb{N}$. In fact, because bounded nets are considered, there does not exist any $\vec{\sigma}$ such that $C \cdot \vec{\sigma} \gneq 0$, and the repetitive sequences are such that $C \cdot \vec{\sigma} = 0$ (i.e. they are *marking repetitive*, $M[\sigma\rangle M$, or *consistent*).

The paper is organized as follows. First we discuss the stochastic interpretation of nets and the possibility of estimating and bounding their performance (Section 2). Then we give a classification of the classes of nets characterized by a unique repeatable firing sequence and we study some of their qualitative properties that can be exploited to derive performance bounds (Section 3). Under these restrictions we derive results that depend only on the mean values and not on the higher moments of the probability distributions of the random variables that describe the timing of the system. In some sense this independence of the probability distribution is a useful generalization of the results, since higher moments of the delays are usually unknown for real cases, and difficult to estimate and assess. Another extension that is possible taking the bounding approach instead of the exact computation, is that we can derive bounds also in the case of *non-ergodic* systems. In the case of Marked Graphs the upper and lower bounds, computed by means of proper Linear Programming problems, are tight, in the sense that it is possible to construct examples of Petri net models with stochastic timings, whose steady-state performances are arbitrarily close to either bound (Section 4). Later we propose a method to construct processes starting from non-safe-marked-graph models, that is an

*This work has been supported by project PA86-0028 of the Spanish Comisión Interministerial de Ciencia y Tecnología (CICYT).

†This work was performed while G. Chiola was visiting the ETSII Zaragoza, with the financial support of the Caja de Ahorros de la Inmaculada de Zaragoza.

extension of a technique originally used by Ramchandani for safe persistent nets [Ram74]. Using this extended method we derive reachable upper bounds for the case of bounded persistent nets, whose computational complexity is polynomial in the size of the net (Section 5). Section 6 presents the case of consistent mono-T-semiflow nets. Conclusions are summarized in Section 7. In the Appendix of the companion paper [CCCS89] the notation and terminology as well as some results are listed for Petri nets.

2 Stochastic interpretation of nets

In the original definition, Petri nets did not include the notion of time. Nevertheless the introduction of timing specification is essential if we want to use this class of models for an evaluation of the performance of distributed systems.

2.1 Timing and firing process

Historically there have been two ways of introducing the concept of time in PN models, namely, associating a time interpretation with either places or transitions; in the latter case transitions have been defined to fire either atomically or in three phases. A more detailed discussion of the timing and firing process can be found in [AMBB*89]. From a modelling point of view, the only effect of these different timing interpretations on the performance evaluation of a model is due to the different implications that the choices have on the resolution of conflicts. Since in the context of this work we are considering subclasses of nets in which either there are no conflicts or their resolution can have no effect on the performance estimates, we would not have to choose one particular interpretation: we could only say that we consider timed PN subclasses. On the other hand, since we are trying to use qualitative results derived from untimed net descriptions, we cannot change the firing mechanism at the level of the net interpretation, since that would invalidate some of the qualitative properties. Hence we must exclude the three-phase firing interpretation. In fact, in a few places we simply speak of marked PNs where we mean timed nets with either timed places or single-phase timed transitions.

2.2 Single versus multiple server semantics: enabling and liveness bounds

Another possible source of confusion in the definition of the timed interpretation of a PN model is the concept of "degree of enabling" of a transition (or reentrance). from the discussion presented in the companion paper [CCCS89] the infinite server semantics appears to be the most general one, and for this reason it will be adopted in this work. Indeed the performance of a model with

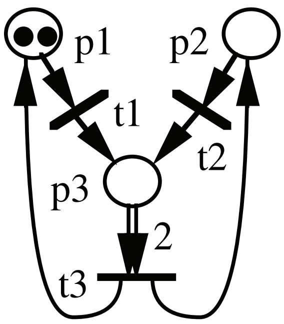

Figure 1: A net with enabling bound greater than the liveness bound for transition t1.

infinite server semantics does depend on the maximum degree of enabling of the transitions; and in particular, the steady-state performance does depend on the maximum degree of enabling of transitions in steady-state, which can be different from the maximum degree of enabling of a transition during its evolution starting from the initial marking, as we shall see in the following. For this reason we recall here the two concepts of degree of enabling of a transition: the enabling bound $E(t)$, and the liveness bound $L(t)$, which allow us to generalize the classical concepts of enabling and liveness of a transition.

DEFINITION 2.1 *Enabling bound.* $\forall t \in T$

$$\begin{aligned} E(t) = \quad & \max \quad k \\ & s.t. \quad \exists M \in R(\mathcal{N}, M_0) \quad : \quad M \geq kPRE[t] \end{aligned}$$

DEFINITION 2.2 *Liveness bound.* $\forall t \in T$

$$\begin{aligned} L(t) = \quad & \max \quad k \\ & s.t. \quad \forall M_1 \in R(\mathcal{N}, M_0) \\ & \qquad \exists M \in R(\mathcal{N}, M_1) \quad : \quad M \geq kPRE[t] \end{aligned}$$

PROPERTY 2.1 *Let $\langle \mathcal{N}, M_0 \rangle$ be a marked PN, then $\forall t \in T \; E(t) \geq L(t)$.*

Now as for any reversible net the reachability graph is strongly connected, the following can be stated.

PROPERTY 2.2 *Let $\langle \mathcal{N}, M_0 \rangle$ be a reversible PN, then $\forall t \in T \; E(t) = L(t)$.*

As a particular case, live Marked Graphs are reversible, so that enabling and liveness bounds are equal in this case. On the other hand, this is not the case for the more general cases that we consider in this paper; indeed the net in Figure 1 with an initial marking of two tokens in p1 and the other two places empty gives an example of a live and bounded PN in which $E(t1) = 2 > L(t1) = 1$ (this net is both—structurally—persistent and mono-T-semiflow, according to the definitions given in Section 3).

A case of strict inequality in Property 2.1 can be interpreted as a generalization of the concept of non-liveness: there exist transitions that "contain potential servers" that are never used in the steady-state; these additional servers might only be used in a transient phase, so that they can die during the evolution of the model. On the other hand it is not difficult to see that the condition $L(t) > 0$ is equivalent to the usual liveness condition for transition t. The problem of the possible difference between enabling and liveness bounds is also related to the ergodicity of the model, as we shall point out later.

The two definitions above refer to behavioural properties that depend on the reachability graph of a PN. Since we are looking for computational techniques at the structural level, we can also introduce structural counterparts of both concepts [CCCS89].

DEFINITION 2.3 *Structural enabling bound:* $\forall t \in T$

$$\begin{aligned} SE(t) = \quad & \max \quad k \\ & s.t. \quad \exists \vec{\sigma} \geq 0, M = M_0 + C\vec{\sigma} \geq kPRE[t] \end{aligned}$$

DEFINITION 2.4 *Structural liveness bound:* $\forall t \in T$

$$\begin{aligned} SL(t) = \quad & \max \quad k \\ & s.t. \quad \forall \vec{\sigma_1} \geq 0, M_1 = M_0 + C\vec{\sigma_1} \geq 0 \\ & \qquad \exists \vec{\sigma} \geq 0, M = M_1 + C\vec{\sigma} \geq kPRE[t] \end{aligned}$$

Note that, the definition of structural enabling bound reduces to the formulation of a Linear Programming problem [Mur83] using matrix C (the incidence matrix of the PN), while the structural liveness bound does not. Now let us look for relations among these enabling bound concepts (the first one is obvious from the trivial implication $M \in R(\mathcal{N}, M_0) \implies M = M_0 + C\vec{\sigma} \wedge \vec{\sigma} \geq 0$).

PROPERTY 2.3 *Let* $\langle \mathcal{N}, M_0 \rangle$ *be a marked PN.*

(i) $\forall t \in T \ SE(t) \geq E(t)$

(ii) $\forall t \in T \ SL(t) \leq SE(t)$

Finally we can point out that, according with Definition 2.4, structural liveness bound is of little practically use because the different domains of the quantifiers do not allow to compare $SL(t)$ and $L(t)$ (i.e. in general we can assert neither $SL(t) \leq L(t)$ nor $SL(t) \geq L(t)$).

2.3 Ergodicity, measurability, bounds

In order to be able to speak about steady-state performance we have to assume that some kind of "average behaviour" can be estimated on the long run of the system we are studying. The usual assumption in this case is that the system models must be *ergodic*, meaning that at the limit when the observation period tends to infinity, the estimates of average values tend (almost surely) to the theoretical expected values of the (usually unknown) probability distributions that characterize the performance indexes of interest.

This assumption is very strong and difficult to verify in general; moreover, it creates problems when we want to include the deterministic case as a special case of a

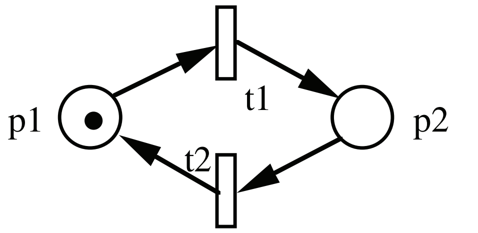

Figure 2: A trivial weakly but non-strongly ergodic deterministic net.

stochastic model, as we will see later on. Thus we introduce also the concept of *weak ergodicity* that allows the estimation of long run performances also in the case of deterministic models.

DEFINITION 2.5 *Ergodicity:*

i) *A (not necessarily stochastic) process* $X(\tau)$ *is said to be weakly ergodic (or measurable in long run) iff the following limit exists:*

$$\overline{X} \stackrel{\text{def}}{=} \lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau X(s)\, ds < \infty$$

ii) *A stochastic process* $X(\tau)$ *is said to be strongly ergodic iff the following condition holds:*

$$\lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau X(s)\, ds = \lim_{\tau \to \infty} E[X(\tau)] < \infty, \ \ a.s.$$

For stochastic Petri nets, ergodicity of the marking and the firing processes can be defined in the following terms:

DEFINITION 2.6 *The marking process of a stochastic marked net is ergodic iff the following limit exists:*

$$\overline{M} \stackrel{\text{def}}{=} \lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau M(s) ds < \vec{\infty}$$

The firing process of a stochastic marked net is ergodic iff the following limit exists:

$$\vec{\sigma}^* \stackrel{\text{def}}{=} \lim_{\tau \to \infty} \frac{\vec{\sigma}(\tau)}{\tau} < \vec{\infty}$$

The (strong) ergodicity concepts [FN85] *are defined in the obvious way taking into consideration Definition 2.5.*

Figure 2 shows a trivial example of a PN in which the marking process is weakly but not strongly ergodic when transitions t1 and t2 are associated with deterministic firing delays θ_1 and θ_2: indeed $E[M(\tau)] = M(\tau)$ is in this case a periodic function of time, so that $\lim_{\tau \to \infty} E[M(\tau)]$ does not exist even if $\overline{M} = (\frac{\theta_2}{\theta_1 + \theta_2}, \frac{\theta_1}{\theta_1 + \theta_2})^T$ in Definition 2.6.

Ergodicity of the marking and of the firing processes are, in general, unrelated properties. For the case of time-invariant (i.e. with transition firing delays independent of time) persistent stochastic Petri nets (that we consider

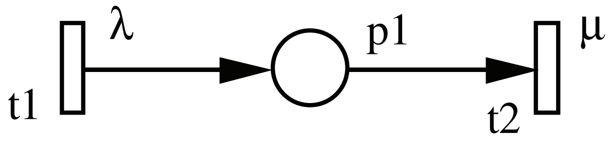

Figure 3: A net with home states but with a non-ergodic marking process.

in Section 5), weak ergodicity of the marking process implies weak ergodicity of the firing process, but not vice-versa. The net in Figure 3 has strongly ergodic firing process but non (even weakly) ergodic marking process if exponentially distributed timing is assumed with rates λ and μ such that $\lambda > \mu$.

The marking process of a bounded model is weakly ergodic if, after a possible transient phase, the system state is always trapped in a unique livelock. In this case $L(t)$ represents the maximum re-entrance of transition t that is guaranteed to be always obtainable; if the model after an initial transient can reach different livelocks, $L(t)$ represents the minimum among the different livelocks of the number of non dead servers for transition t in each one. From this point of view, it makes sense to *compute upper and lower bounds on transition firing frequences also in the case of non weakly ergodic models* having more than one livelock, even if the "true" average value does not exist (i.e. the limit for $\tau \to \infty$ is not unique).

THEOREM 2.1 *If a marked Petri net has a unique repetitive firing count vector, then its firing processes is weakly ergodic.*

PROOF: If the net has a unique repetitive firing count vector $\vec{\sigma_R}$, then by Theorem 3.9 in [Sil87]

$$\vec{\sigma}(\tau) = \frac{\tau}{\Gamma}\vec{\sigma_R} + b(\tau)$$

where $b(\tau)$ is an almost surely bounded vector, and Γ is the expected value of the random duration of one execution of $\vec{\sigma_R}$. Then the firing process is weakly ergodic since:

$$\lim_{\tau\to\infty} \frac{\vec{\sigma}(\tau)}{\tau} = \frac{\vec{\sigma_R}}{\Gamma} = \vec{\sigma^*}$$

Q.E.D.

The subclass of bounded and persistent nets are guaranteed to be weakly ergodic, in the sense that they may have at most one livelock (see Theorem 3.2).

DEFINITION 2.7 *A stochastic Petri net is said to be semi-Markovian iff the related marking process is a semi-Markov process.*

THEOREM 2.2 *If a semi-Markovian marked Petri net has a home state and the mean value of the marking remains bounded ($\sup_\tau E[M(\tau)] < \infty$, $\tau \in I\!R^+$) then either the marking process is strongly ergodic or the underlying Markov process is recurrent null.*

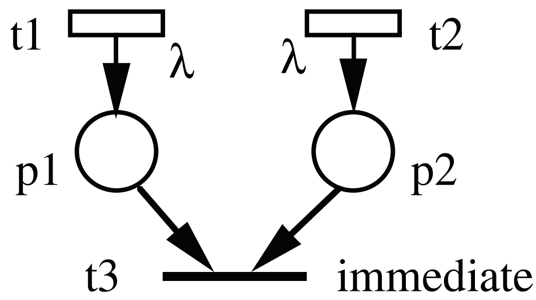

Figure 4: An exponential net with a home space but with a non strongly ergodic marking process.

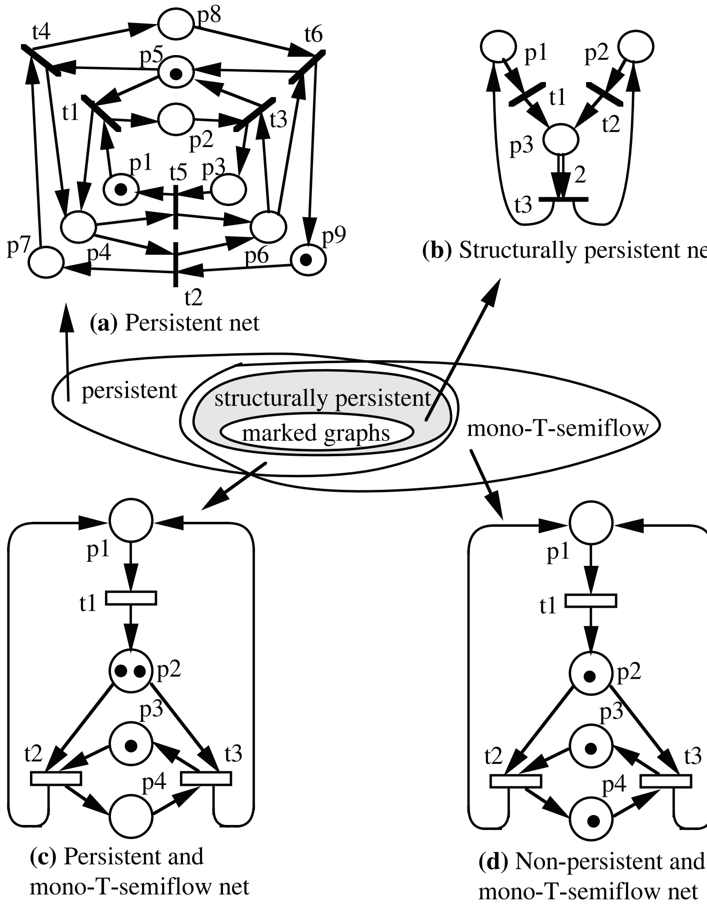

Figure 5: Relation among live and bounded net subclasses having a unique consistent firing count vector.

PROOF: The home space (the set of home states of the net) is the unique recurrent class of the underlying Markov process. If this class is positive recurrent, then the marking process is strongly ergodic.
Q.E.D.

The complete separation of the null recurrent case is not established. This separation is difficult from a theoretical point of view, and it has little practical interest. Figure 4 shows a Markovian net with a home space but with a non-ergodic marking process.

In Figure 3 a (unbounded) net is shown with home state but non ergodic marking process if the mean value of the random variable associated with transition t_1 is less than the one associated with t_2. On the other hand, nets can have bounded marking mean values and be non ergodic because of the presence of more than one livelock.

The reverse of the previous theorem is not true. If the marking process of a semi-Markovian net is strongly ergodic then

$$\lim_{\tau\to\infty}\frac{1}{\tau}\int_0^\tau M(s)ds = \lim_{\tau\to\infty} E[M(\tau)] < \infty, \text{ a.s.}$$

so the mean value of the marking remains bounded. But the existence of a home state cannot be assured. For the net in Figure 8, an exponential distribution timing can be associated with transitions such that the related marking process is ergodic (e.g. taking the same value for the rates of all transitions) and the net has no home state.

An important particular case of semi-Markovian Petri nets is that of *Coxian* distribution timing of transitions. The Coxian family is generated from exponential distributions by convolutions (*generalized Erlangs*) and mixtures (*hyperexponential* distributions) (see [Cox55], and [Har85]). They are characterized by having rational Laplace transform. Petri nets with Coxian distribution timing are Markovian nets (i.e. the related marking process is Markov). The interest of this family of distributions is that any distribution function can be approximated with a Coxian, preserving mean and variance [GP87].

From Theorem 2.2 trivially follows that

COROLLARY 2.1 *If a semi-Markovian bounded marked net has a home state then its marking process is strongly ergodic.*

Intuitively, the previous result can be interpreted in the following way: The possibly different behaviours of the marking process between successive arrivals to a home state are statistically equivalent (a renewal process can be defined in terms of the number of arrivals to that home state). The length of time between these arrivals has finite mean (because the net is bounded). Then, for a given trajectory, time averages lead to the same value as ensemble averages (each behaviour between two arrivals can be seen as a new execution of the process) at the renewal times.

3 Live and bounded nets with a unique repetitive firing count vector

Nets with a unique repetitive firing count vector can be obtained from two non disjoint subclasses of live and bounded nets, named *persistent* and *mono-T-semiflow nets*. Persistent nets are behaviourally defined, while mono-T-semiflow are structurally characterized. As a particular case, *structurally persistent nets* and *marked graphs* belong to the intersection of these two classes, thus possessing the good properties of both. This section is devoted to the introduction of these nets subclasses and to the presentation of some of their basic properties relevant from the performance point of view. Figure 5 provides an overall picture of the relations among the net subclasses considered in this work.

3.1 Persistent nets

DEFINITION 3.1 *A marked Petri net $\langle \mathcal{N}, M_0\rangle$ is said to be persistent iff for all reachable marking M and for all different transitions, t_1 and t_2, enabled in M, the sequence $t_1 t_2$ is fireable from M.*

Persistent nets are effectively conflict-free nets. As an example look at the net in Figure 5a. This net has structural conflicts but for the initial marking $M_0 = (1,0,0,0,1,0,0,0,1)^T$ does not reach any state in which a decision must be taken.

PROPERTY 3.1 *Let $\langle \mathcal{N}, M_0\rangle$ be a persistent net, and $M_0' \leq M_0$. Then $\langle \mathcal{N}, M_0'\rangle$ is persistent.*

PROOF: Let M' be a reachable marking from M_0', i.e. such that there exists a sequence σ such that $M_0'[\sigma\rangle M'$. Since $M_0' \leq M_0$, σ can be fired also when starting from M_0 instead of M_0'. Moreover, if $M_0[\sigma\rangle M$ then $M \geq M'$. Since $\langle \mathcal{N}, M_0\rangle$ is persistent there is no effective conflict in M and therefore there is no conflict in M'.
Q.E.D.

Let us now introduce some concepts (see the Appendix of [CCCS89] for their formal definition) and results that will lead to the conclusion that *the stochastic marking process associated with a bounded persistent net is weakly ergodic*. The first concept we introduce is that of *directedness*; this means that any two reachable markings have at least one common successor marking. A second concept is that of *home states*, markings which can be reached from any other reachable marking. The set of home states of a marked net forms a unique livelock. Finally, a marked net is *reversible* if the initial marking can always be recovered from any reachable marking (i.e., M_0 is a home state).

LEMMA 3.1 [Bra83] *All persistent marked nets have the directedness property.*

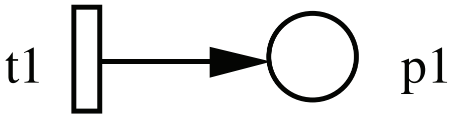

Figure 6: An unbounded live and persistent net having the directedness property but no home states.

LEMMA 3.2 [BV84] *For bounded marked nets, directedness and the existence of a home state are equivalent properties.*

Figure 6 illustrates an example that shows that the above lemma does not hold for unbounded nets: the net depicted there is unbounded, has the directedness property and has no home states.

THEOREM 3.1 *Live bounded persistent connected nets without implicit places have a unique repetitive firing count vector (i.e. $\exists \vec{\sigma}_R$ such that $\forall M \in R(\mathcal{N}, M_0)$ if $M[\sigma\rangle M$ then $\vec{\sigma} = k.\vec{\sigma}_R$ with $k \in \mathbb{N}$).*

PROOF: Consider a live, bounded, persistent, connected net $\langle \mathcal{N}, M_0 \rangle$. Bounded persistent nets have an home space, so that let M be a home state of the net. Since the net is live, there exist at least one firing sequence σ such that $M[\sigma\rangle M$ and $\vec{\sigma} \geq \vec{1}$. Now assume that there exist two different repetitive firing count vectors $\vec{\sigma_1}$ and $\vec{\sigma_2}$ such that $M[\sigma_1\rangle M$ and $M[\sigma_2\rangle M$ and $\vec{\sigma_1} + \vec{\sigma_2} \geq \vec{1}$. Then, there must exist three transitions t_i, t_j, and t_k such that $t_i \in ||\vec{\sigma_1}||$ and $t_i \not\in ||\vec{\sigma_2}||$ and $t_j \in ||\vec{\sigma_2}||$ and $t_j \not\in ||\vec{\sigma_1}||$ and $t_k \in ||\vec{\sigma_1}||$ and $t_k \in ||\vec{\sigma_2}||$. Moreover, since the net is connected and each firing count vector is a consistent component, there must be a structural conflict between the two transitions t_i and t_j, i.e., $\exists p \in t_k{}^\bullet$ such that $p \in {}^\bullet t_i \cap {}^\bullet t_j$. Since the net is persistent, the structural conflict between t_i and t_j cannot be effective, and the two sequences σ_1 and σ_2 are firable independently one of the other, so that the shared place p must be behaviourally implicit.
Q.E.D.

Since live bounded persistent connected nets have a unique repetitive firing count vector we can obtain the following particularization of Theorem 3.9 in [Sil87].

THEOREM 3.2 *For live bounded connected marked nets without implicit places, persistency implies weak ergodicity of the firing process.*

PROOF: By Theorem 3.1, for live bounded connected marked nets without implicit places, persistency implies the existence of a unique repetitive firing count vector. Then by Theorem 2.1 for these nets, persistency implies weak ergodicity of the firing process.
Q.E.D.

In order to study the steady-state performances of a stochastic net, only *recurrent* markings are relevant (i.e. *transient* markings do not affect the computation). Even if bounded persistent nets are ergodic this does not mean that there exist no transient markings. The net in Figure 1 is structurally persistent, live, and 2–bounded for $M_0 = (2, 0, 0)^T$, but M_0 is a transient state (i.e. it is not a home state). Nets without transient markings are said to be *reversible*.

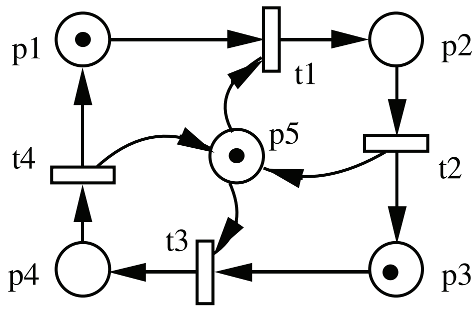

Figure 7: A live and bounded mono-T-semiflow net.

From the definition of reversible nets (Definition A.3 in [CCCS89]) and Corollary 2.1, the following can be easily concluded:

COROLLARY 3.1 *Semi-Markovian stochastic reversible and bounded marked nets have ergodic marking process.*

3.2 Mono-T-semiflow nets

Let us introduce another class of nets with a unique repetitive firing count vector which are structurally characterized.

DEFINITION 3.2 *A structurally bounded Petri net $\mathcal{N}$ is called mono-T-semiflow iff there exists a unique minimal T-semiflow that contains all transitions.*

In a mono-T-semiflow net conflicts may be reached, and so different behaviours can occur. However, from the steady-state performances point of view, these decisions lead us to a unique result, provided that the net is live and bounded (all different behaviours let the same set of transitions, characterized by the only T-semiflow of the net, fire, perhaps in a different order). For example, $\sigma_a = t_1 t_2 t_3 t_4$ and $\sigma_b = t_3 t_4 t_1 t_2$ are possible sequences in the net of Figure 7, both fireable from M_0. Even if the performance is equal for any conflict resolution policy, from the functional point of view the results can be very different (imagine $t_1 t_2$ and $t_3 t_4$ be two non commutative operations).

The following result is the analogous of Theorem 3.1 for mono-T-semiflow nets.

PROPERTY 3.2 *Let $\langle \mathcal{N}, M_0 \rangle$ be a live and bounded mono-T-semiflow net. For any firing sequence σ applicable in $\langle \mathcal{N}, M_0 \rangle$ we can write: $\vec{\sigma} = r.X + b$, where $X > 0$ is the minimal T-semiflow, $r \in \mathbb{N}$, and $b \geq 0$ is a bounded vector such that $C \cdot b \gneq 0$ is impossible.*

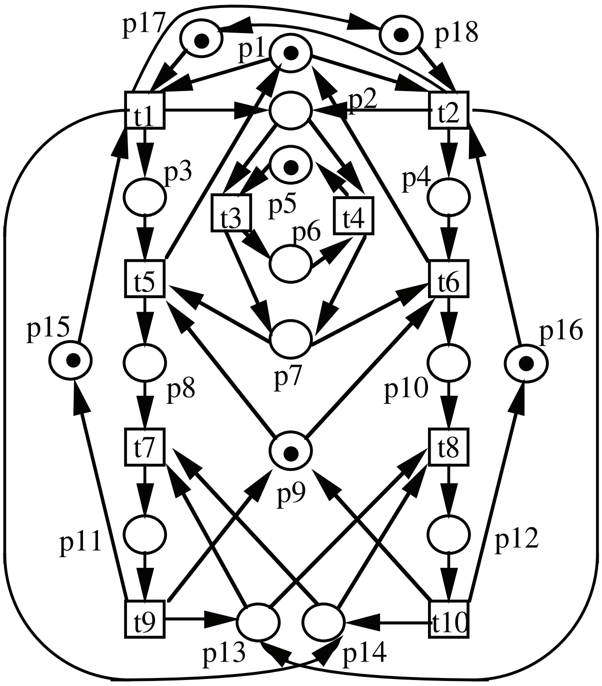

Figure 8: A live bounded mono-T-semiflow net without home states.

COROLLARY 3.2 *For mono-T-semiflow bounded nets, deadlock-freeness and liveness are equivalent properties.*

The next property allows a polynomial complexity characterization of mono-T-semiflow nets:

PROPERTY 3.3 *Let C be the incidence matrix of a consistent mono-T-semiflow net. Then $rank(C) = m-1$.*

PROOF: If a net is mono-T-semiflow then the dimension of the space of right annihilators of C is great than or equal to 1, thus $rank(C) \leq m-1$. By contradiction, suppose X and b are two linearly independent right annihilators of C. By consistency X can be chosen positive, and $b \in Z^n$. Then two independent and positive right annihilators of C can be constructed: $X_1 = \alpha X - \beta b$, $X_2 = \delta X + \gamma b$ (with proper values for α, β, γ, and δ), thus X_1, X_2 are independent T-semiflows and this is against the hypothesis of mono-T-semiflow.
Q.E.D.

We have seen that bounded and persistent nets have an ergodic marking behaviour. Now, what can be said about mono-T-semiflow nets? Unfortunately the net in Figure 8 [BV84] shows a negative result: it is live and structurally bounded (and therefore consistent and conservative), mono-T-semiflow, but has no home states since it generates two livelocks, and thus it may lead to non ergodic marking processes. In fact, if there exists a positive probability of reaching two different livelocks, the ergodicity is impossible.

3.3 Structurally persistent nets and marked graphs

Persistency is a behavioural property, i.e. the same net structure with a different initial marking can give non persistent behaviour. For example, the net in Figure 5c is persistent. On the same structure but with $M_0 = (0,1,1,1)^T$ the net is not persistent (Figure 5d). In both cases the net is live and bounded, and its unique repetitive firing count vector is given by $X = (2,1,1)^T$. Let us introduce a subclass of persistent nets such that the persistency is inherent to the structure.

DEFINITION 3.3 *A Petri net $\mathcal{N}$ is said to be structurally persistent iff $\langle \mathcal{N}, M_0 \rangle$ is persistent for all finite initial marking M_0.*

PROPERTY 3.4 [LR78] *A net is structurally persistent iff it does not exist any structural conflict in it.*

The live and bounded structurally persistent net in Figure 5b shows that they may have transient states. Now it is interesting to point out a very well known subclass of these nets for which there exists no transient markings, provided the liveness for M_0.

DEFINITION 3.4 [CHEP71] *Marked Graphs are ordinary Petri nets (pre and post incidence functions taking values in $\{0,1\}$) such that $|{}^\bullet p| = |p^\bullet| = 1, \forall p \in P$.*

PROPERTY 3.5 *Marked graphs are structurally persistent nets. The reverse in not true (for example the net in Figure 5b is not a marked graph).*

PROPERTY 3.6 *Classification*

i) *Structurally persistent live and bounded nets are mono-T-semiflow. The reverse is not true (Figure 5c).*

ii) *Marked graphs are consistent nets, and the unique minimal T-semiflow is $\vec{1}$.*

By properties 3.6 and 3.3, as marked graphs are consistent nets, the rank of their incidence matrix is m$-$1.

THEOREM 3.3 [Mur77] *Let $\langle \mathcal{N}, M_0 \rangle$ be a live (possibly unbounded) marked graph. The two following statements are equivalent:*

i) *$M \in R(\mathcal{N}, M_0)$, i.e. M is reachable from M_0.*

ii) *$B_f \cdot M = B_f \cdot M_0$, with B_f the fundamental circuit matrix of the graph (i.e. its row vectors are a basis of the left annihilators of C), and $M \geq 0$.*

According to the above theorem $M \in R(\mathcal{N}, M_0) \iff M_0 \in R(\mathcal{N}, M)$. In other words:

PROPERTY 3.7 *Live marked graphs are reversible.*

PROPERTY 3.8 *Let $\mathcal{N}$ be a marked graph.*

i) *$\mathcal{N}$ is structurally bounded (i.e. $\langle \mathcal{N}, M_0 \rangle$ is bounded $\forall M_0$) iff it is strongly connected.*

ii) *Let $\langle \mathcal{N}, M_0 \rangle$ be live. Then $\langle \mathcal{N}, M_0 \rangle$ is bounded iff $\mathcal{N}$ is structurally bounded.*

According to the above properties, strong connectivity and boundedness have equivalent meaning for live MGs. From Properties A.2, A.3 in [CCCS89], and Corollary 3.1, the following is thus obvious:

COROLLARY 3.3
Semi-Markovian strongly connected live marked graphs have ergodic marking process.

Finally, an interesting property (that shall be used later in Section 4.3) of live MGs, that allows an efficient computation of liveness bounds, is the following:

PROPERTY 3.9 *Let $\langle \mathcal{N}, M_0 \rangle$ be a live MG, then $\forall t \in T$ $SE(t) = E(t) = SL(t) = L(t)$.*

In other words, for marked graphs the behavioural concepts always collapse into the structural ones. In fact, according to Theorem A.1 in [CCCS89], for any MG $M \in R(M_0)$ iff $M = M_0 + C\vec{\sigma} \wedge \vec{\sigma} \geq 0$.

In case of non-strongly connected MGs it is possible to obtain $SE(t) = \infty$ for some transition t; this creates no harm: because of the assumption of an infinite server semantics, it only implies that the timing of that transition does not affect the steady-state performance of the model.

4 Bounds for stochastic strongly connected marked graphs

In this Section, performance bounds for strongly connected (and thus structurally bounded) MGs are recalled. Strong connectivity of a graph is a well known problem of polynomial time complexity.

4.1 Upper bound for the steady state throughput

Let us take into account just the first moments of the probability density functions associated with transitions. In the following, let θ_i be the mean value of the random variable associated with the firing of transition t_i, and D the diagonal matrix with elements θ_i, $i = 1, ..., m$.

The limit expected firing count vector per time unit is

$$\vec{\sigma}^* = \lim_{\tau \to \infty} \frac{\vec{\sigma}(\tau)}{\tau} \qquad (1)$$

and the mean time between two consecutive firings of a selected transition, t_i,

$$\Gamma_i = \frac{1}{\vec{\sigma}_i^*} \qquad (2)$$

Then the components of $Pre \cdot D \cdot \vec{\sigma}^* \cdot \Gamma_i$ (where $\vec{\sigma}^*$ is normalized by Γ_i) represent the product of the number of tokens reserved for firing the transitions and the mean length of time that these tokens reside in each place between two consecutive firings of t_i.

For nets with a unique repetitive firing count vector, $Pre \cdot D \cdot \vec{\sigma}^* \cdot \Gamma_i = Pre \cdot D \cdot X$, where $X_i = 1$. If the net is structurally bounded mono-T-semiflow then X is its minimal T-semiflow. If the net is persistent (Figure 5a) then X is a T-semiflow (possibly non minimal).

Let $\overline{M}$ be the limit vector of the average number of tokens in each place (i.e. $\overline{M} = \lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau M(s)ds$). Then, provided that the previous limit exists, $\overline{M} \cdot \Gamma_i$ is the vector of products of the mean number of tokens and the length of one cycle and we have:

$$\overline{M} \cdot \Gamma_i \geq Pre \cdot D \cdot X \qquad (3)$$

From this inequality, the minimum cycle time associated with transition t_i, Γ_i^{min}, can be derived. We take into account that Γ_i^{min} must be such that inequality (3) holds and for some place p_j the equality is reached:

$$\Gamma_i^{min} = \frac{Pre(j) \cdot D \cdot X}{\overline{M}(j)} \qquad (4)$$

Since the vector $\overline{M}$ is unknown, (4) cannot be solved. Making the product with a P-semiflow Y for any reachable marking M:

$$Y^T \cdot M_0 = Y^T \cdot M = Y^T \cdot \overline{M} \qquad (5)$$

Now, from (3) and (5):

$$Y^T \cdot M_0 \cdot \Gamma_i \geq Y^T \cdot Pre \cdot D \cdot X \qquad (6)$$

And the minimum cycle time in steady state is:

$$\Gamma_i^{min} = \max_{Y \in \{\text{P-semiflow}\}} \frac{Y^T \cdot Pre \cdot D \cdot X}{Y^T \cdot M_0} \qquad (7)$$

Of course, an upper bound for the throughput of t_i is $\frac{1}{\Gamma_i^{min}}$.

Let us formulate the previous lower bound for the cycle time in terms of a particular class of optimization problems that we introduce now.

DEFINITION 4.1 *A fractional programming problem is an optimization problem of the form* [Mur83]*:*

$$\begin{array}{lll} \max & f(x) = & \dfrac{c^T \cdot x + \alpha}{d^T \cdot x + \beta} \\ \textit{subject to} & & A \cdot x = b; \quad x \geq 0 \\ & & A \in \mathbb{R}^{v \times u}; \quad b \in \mathbb{R}^v \\ & & c, d \in \mathbb{R}^u; \quad \alpha, \beta \in \mathbb{R} \end{array}$$

A fractional programming problem is said to be homogeneous if coefficients α and β are equal to zero.

THEOREM 4.1 *For any persistent or mono-T-semiflow net, the minimum cycle time associated with transition t_i can be computed by the following homogeneous fractional programming problem:*

$$\begin{aligned} \Gamma_i^{min} = \quad & \max \quad \frac{Y^T \cdot Pre \cdot D \cdot X}{Y^T \cdot M_0} \\ & s.t. \quad Y^T \cdot C = 0 \qquad\qquad (FPP1) \\ & \qquad Y \geq 0, \quad \vec{1}^T \cdot Y > 0 \end{aligned}$$

PROOF: Just notice that the optimum can be always reached with an elementary P-semiflow. If Y^* gives us the optimal solution, Z^*, and it is not an elementary P-semiflow, then

$$Y^* = \sum_{i=1}^{k} (a_i Y_i) \tag{8}$$

with Y_i elementary P-semiflows and $a_i > 0, \sum_{i=1}^{k} a_i = 1$ (i.e. $Z^* = \sum_{i=1}^{k} (a_i Y_i^T \cdot Pre \cdot D \cdot X)$). Then $a_i = \frac{1}{k}$ $\forall i$, all $Y_i^T \cdot Pre \cdot D \cdot X$ must be equal (if not, the maximum of them would give a larger value than Z^*) and all Y_i are optimal solutions of the problem.
Q.E.D.

4.2 Upper bounds for strongly connected marked graphs

Because marked graphs have a unique T-semiflow $X = \vec{1}$ (Property 3.6), writing $\theta = D \cdot \vec{1}$ the following can be directly stated:

COROLLARY 4.1 *For strongly connected marked graphs, the minimum cycle time can be obtained by solving the following fractional programming problem:*

$$\begin{aligned} \Gamma^{min} = \quad & \max \quad \frac{Y^T \cdot Pre \cdot \theta}{Y^T \cdot M_0} \\ & s.t. \quad Y^T \cdot C = 0 \qquad\qquad (FPP2) \\ & \qquad Y \geq 0, \quad \vec{1}^T \cdot Y > 0 \end{aligned}$$

THEOREM 4.2 [CCCS89] *For live strongly connected marked graphs, the minimum cycle time can be obtained by solving the following linear programming problem:*

$$\begin{aligned} \Gamma^{min} = \quad & \max \quad Y^T \cdot Pre \cdot \theta \\ & s.t. \quad Y^T \cdot C = 0, Y^T \cdot M_0 = 1 \qquad (LPP1) \\ & \qquad Y \geq 0 \end{aligned}$$

Theorem 4.2 shows that the problem of finding an upper bound for the steady-state throughput in a strongly connected stochastic marked graph can be solved looking at the cycle times associated with each P-semiflow (cycles for marked graphs) of the net, considered in isolation. These cycle times can be computed making the summation of the average firing times of all the transitions involved in the P-semiflow, and dividing by the number of tokens present in it.

The above bound is the same that has been obtained for strongly connected deterministic marked graphs by other authors (see for example [Ram74], [RH80], [Mur85]), but here it is considered in a practical LPP form. For these nets, the reachability of this bound has been shown ([Ram74], [RH80]). Since deterministic timing is just a particular case of stochastic timing, the reachability of the bound is assured for our proposes. Even more, the next result shows that the previous bound cannot be improved only on the base of the knowledge of the coefficients of variation for the transition firing times.

THEOREM 4.3 [CCCS89] *For strongly connected marked graphs with arbitrary values of mean and variance for transition firing times, the bound for the cycle time obtained from (FPP2) cannot be improved.*

4.3 Lower bounds for strongly connected marked graphs

A trivial lower bound in steady-state performance for a live PN with a unique repetitive firing count vector is of course given by the sum of the firing times of all the transitions weighted by the firing count vector itself. Since the net is live all transition must be fireable, and the sum of all firing times multiplied by the number of occurrences of each transition in the (unique) average cycle of the model corresponds to any complete sequentialization of all the activities represented in the model. This lower bound is always reached in an MG consisting of a single loop of transitions and containing a single token in one of the places, independently of the higher moments of the PDFs (this observation can be trivially confirmed by the computation of the upper bound, which in this case gives the same value).

This trivial lower bound has been improved in [CCCS89], based on the knowledge of the liveness bound $L(t)$ for all transitions t of the MG.

THEOREM 4.4 [CCCS89] *For any live and bounded MG with a specification of the mean firing times θ_j for each $t_j \in T$ it is not possible to assign PDFs to the transition firing times such that the average cycle time is greater than*

$$\Gamma_{max} = \sum_j \frac{\theta_j}{L(t_j)}$$

independently of the topology of the net (and thus independently of the potential maximum degree of parallelism intrinsic in the MG).

The lower bound in performance given by the computation of Γ_{max} as defined in Theorem 4.4 has been shown to be reachable for any MG topology and for some assignement of PDF to the firing delay of transitions in [CCCS89]. First of all we recall that in the case of live MGs the liveness bound equals the structural enabling bound for each transition (Property 3.9); thus the problem of the determination of the structural enabling bound

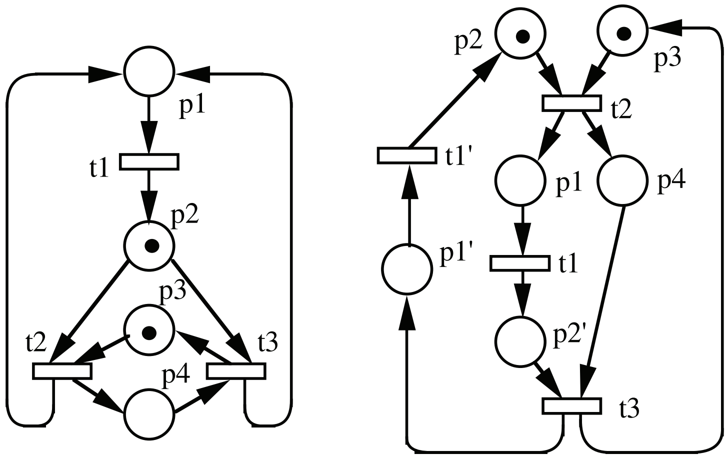

Figure 9: Equivalent marked graph for safe persistent net.

can be characterized in terms of a Linear Programming problem, which is known to be solvable in polynomial time.

For any transition $t \in T$, the computation of the structural enabling bound $SE(t)$ can be formulated in terms of the following LPP:

$$\begin{array}{rll} SE(t) = & \min & Y^T M_0 \\ & \text{s.t.} & Y^T C = 0 \\ & & Y^T PRE[t] = 1 \\ & & Y^T \geq 0 \end{array}$$

Because of the minimization requirement, the optimum of the objective function is always achieved with elementary P-semiflows Y. In case of MGs, these elementary P-semiflows can only be elementary cycles, so that we can give the following interpretation of the LPP in net terms: the liveness bound for a transition t of a strongly connected MG is given by the minimum number of tokens contained in any cycle of places containing transition t. In a non-strongly connected MG there can be no such cycle, so that this number can be infinite.

5 Throughput bounds for bounded persistent nets

As we said in Section 3, persistent nets are behaviourally defined. This means that a previous behavioural analysis to assure the persistency of the net must be made before computing bounds for the performances in steady-state. Few results are known in the literature related with bounds for the performances of stochastic bounded persistent nets. A partial result is presented in [Ram74] for safe and persistent nets with deterministic timing. For these nets a behaviourally equivalent safe marked graph can be built. The method consists in drawing the initially marked places and enabled transitions. After that, firing all transitions and drawing the output places and repeating the procedure until a marking in the process is re-found (see Figure 9). Then, the methods explained in Sections 4 and 5 can be applied for computing the bounds for this marked graph and so for the steady-state performances of the initial persistent net. Unfortunately, this analysis is not possible for bounded (non safe) nets when non deterministic timing is considered.

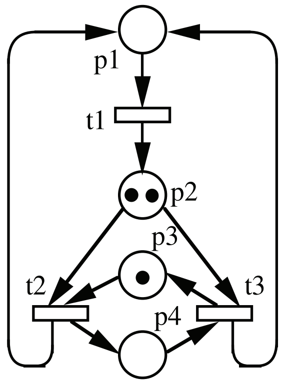

Figure 10: Bound (LPP3) non reachable.

Let us now introduce some general results useful for computing bounds for the performance of bounded persistent nets. Later we improve some of these results.

5.1 General results for bounded persistent nets

For bounded persistent nets, ergodicity of the marking and the firing processes is assured for both single and three phases firing interpretation (Theorem 3.2). Then for these nets it makes sense to speak about the unique steady-state behaviour and to compute bounds for the performance of this steady-state.

5.1.1 Upper bound for the throughput

Let us consider live, bounded and persistent nets without implicit places. According to Theorem 3.1 consistent firing count vectors are proportional to $\vec{\sigma}_R$.

The next linear programming problem follows from the general problem (FPP1) in Section 4. It can be considered for computing a lower bound of the steady-state cycle time of a selected transition t_i.

$$\begin{array}{rll} \Gamma_i^{min} = & \max & Y^T \cdot Pre \cdot D \cdot X \\ & \text{s.t.} & Y^T \cdot C = 0, Y^T \cdot M_0 = 1 \qquad (LPP3) \\ & & Y \geq 0 \end{array}$$

where D is the diagonal matrix of mean values of the random variables associated with transitions and $X = k\vec{\sigma}_R$ is a T-semiflow (non minimal if there exist more than one) with $X(t_i) = 1$ (the unique repetitive firing count vector). The optimal value of the previous problem is a non reachable bound in general (i.e. there exist net models such that no stochastic interpretation allows to reach the computed bound, Γ_i^{min}). Let us consider, for example, the net in Figure 10. Selecting the transition t_2, the vector X for (LPP3) is $X = (2, 1, 1)^T$ and the obtained bound is $\Gamma_2^{min} = \max(\theta_2 + \theta_3, 2\theta_1)$. Now, considering deterministic timing for all transitions with $\theta_1 = 2, \theta_2 = 60, \theta_3 = 1$, the obtained bound is $\Gamma_2^{min} = 61$ while the actual cycle time for transition t_2 is bigger because of the sequence θ_1, θ_2 which takes 62 units of time.

For safe persistent nets, the previous bound can be always reached. In fact the same bound would be obtained by deriving the equivalent Marked Graph (according to [Ram74]) and computing the bound for it (using the results in Section 4).

Theorem 3.1 states that for connected bounded persistent nets without implicit places, liveness implies the existence of a unique consistent repetitive firing count vector. Now we prove a reverse result:

THEOREM 5.1 *Persistent nets with a consistent repetitive firing count vector are live.*

PROOF: Because of persistency if a transition is enabled it will be fired. Then liveness is assured if the unique repetitive firing count vector is consistent (all transitions belong to the repetitive sequence).
Q.E.D.

Now from Theorems 3.1 and 5.1 the next result follows:

COROLLARY 5.1 *For connected bounded persistent nets without implicit places, liveness and the existence of a consistent repetitive firing count vector are equivalent.*

Even if the cycle time bound obtained from (LPP3) can be non reachable, now we can prove that it is finite if and only if the actual cycle time is finite.

THEOREM 5.2 *For bounded persistent nets with a unique consistent repetitive firing count vector, the lower bound for the cycle time obtained from (LPP3) is finite iff the actual cycle time is finite.*

PROOF: ($\Leftarrow$) Obvious (the value obtained from (LPP3) is a lower bound)

($\Rightarrow$) From Theorem 5.1 the nets that we consider are live so the cycle time in steady-state is finite.
Q.E.D.

The previous result is not true for the general case. We see in Section 6 that for mono-T-semiflow nets the actual cycle time can be infinite (so that the net can be non live) while the lower bound obtained from (LPP3) can be finite.

5.1.2 Lower bound for the throughput

The trivial lower bound presented in Section 5 for the throughput of Marked Graphs can be applied now in the following way: just making the sum of the firing times of the transitions belonging to the unique repetitive firing sequence weighted by the firing count vector itself.

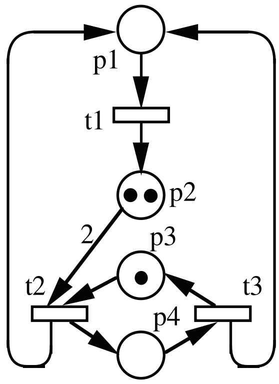

Figure 11: Non trivial lower bound non reachable.

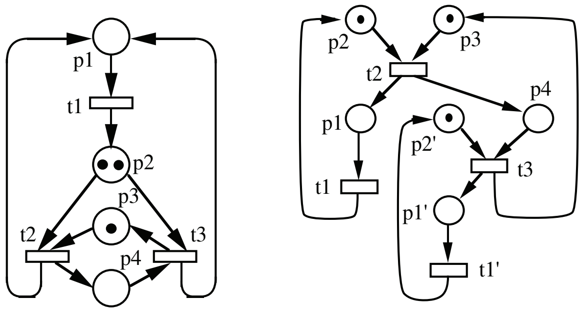

Figure 12: Equivalent marked graph for deterministic timing.

This lower bound is, in general, non reachable. Nevertheless, for safe persistent nets the bound is reachable as can be seen just deriving the behaviourally equivalent safe marked graph as in [Ram74]. Unfortunately, the non trivial lower bound for marked graphs (dividing by the liveness degree) cannot be applied to persistent nets. For example, for the net in Figure 11 the steady-state cycle time is less than the value obtained applying this technique (so that the throughput is higher).

5.2 Improving the upper bound

Let us now describe a method for computing a reachable upper bound for the steady-state throughput for bounded persistent nets. Deterministic timing will lead to the maximum possible performance for a given net and average firing associated to transitions. If we consider only deterministic timing, a behaviourally equivalent marked graph can be derived in a way analogous to that proposed in [Ram74]. By splitting the places in such a way that they represent conditions for the enabling of transitions (i.e. the obtained marked graph is safe). Developing the process until cyclicity appears (i.e. a marking is repeated for the first time) and identifying those instances of places that must be superposed taking into account that the deterministic assumption restricts the behaviour of the net (see the example in Figure 12). Considering general timing distributions, the original net and the derived marked graph are not behaviourally equivalent. In fact, the steady-state throughput for the marked graph is less than or equal to the one of the original net. Nevertheless for deterministic timing the equality holds and this provides the following method for computing a reachable bound for the throughput:

Consider a bounded persistent net with general distribution timing.

Step 1. Develop the cyclic process for the deterministic case (behaviourally equivalent marked graph for the deterministic timing)

Step 2. Compute the upper bound for the steady-state throughput of the marked graph (Section 4.2)

Step 3. The bound computed in Step 2 is a reachable bound for the steady state throughput of the original net.

The bound is reachable because the maximum throughput for the net is always obtained in the deterministic case and under this condition the throughput of the cyclic process and of the original net are equal.

6 Throughput bounds for consistent mono-T-semiflow nets

Let us now consider consistent mono-T-semiflow nets and give bounds for their steady-state throughput. Since mono-T-semiflow nets are structurally characterized, they can be recognized without a previous behavioural analysis. According to the results in Section 3.2, they can be recognized by computing the T-semiflow X in polynomial time. Unfortunately, the existence of a unique consistent minimal T-semiflow does not assure the ergodicity of the marking process (see Section 3) (in general, consistent mono-T-semiflow nets may have no home state, so ergodicity is not assured).

Even in the case in which ergodicity is not assured the problem of computing bounds for the throughput makes sense. The values that we compute in this section are bounds for all possible steady-state behaviours of the net.

The problem (LPP3) for computing an upper bound for the steady-state throughput of a net with a unique repetitive firing count vector (Section 4) can be used. Nevertheless, this bound is, in general non reachable. Moreover, a mono-T-semiflow net can be non live and the obtained lower bound for the cycle time be finite. In other words:

PROPERTY 6.1 *For consistent mono-T-semiflow nets, liveness is not characterized by the finiteness of the lower bound of the cycle time computed by means of LPP3.*

This can be easily checked by considering the net in Figure 13 (which is non-live, so that the actual steady-state cycle time is infinite, even if the obtained bound is finite).

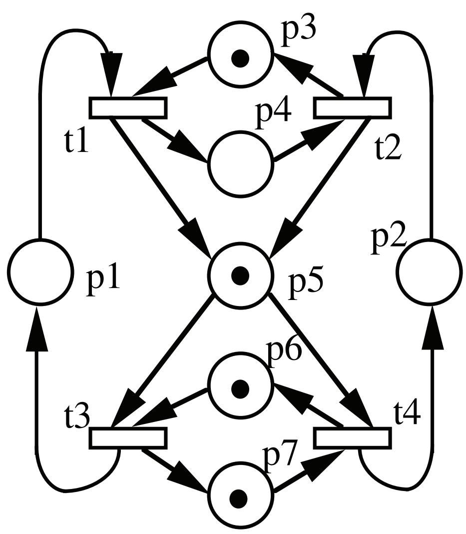

Figure 13: Non live mono-T-semiflow net.

In conclusion, only the trivial lower bound for the throughput of consistent mono-T-semiflow nets (the sum of all transition firing times) can be computed, provided that the net is live.

7 Conclusions

In this paper we have addressed the problem of computing upper and lower bounds for the throughput of transitions in Petri net models having a unique repetitive firing count vector. The results presented here represent an extension of those described in [CCCS89] for the case of bounded Marked Graphs. The upper bound in case of Persistent nets is a generalization of that obtained for Marked Graphs, and has been shown to be reachable. The technique proposed for the derivation of processes for non-safe persistent nets is an extension of a method originally proposed by Ramchandani for safe persistent nets, and that was not directly applicable.

For what concerns the lower bound on throughput, only the trivial bound computed as the inverse of the sum of all transition firing delays can be generalized from the Marked Graph case, and this can be too pessimistic in case of Persistent or mono-T-semiflow nets.

In any case both the upper and lower bounds are independent of any assumption on the probability distribution of the delay associated with transitions, and their value can be computed based on the knowledge of the averages. This represents a generalization with respect to the usual assumptions needed for the exact performance evaluation of a Petri net model. A second generalization coming from the choice of computing bounds instead of actual values is that the analysis of non-(weakly)ergodic models still makes sense.

Besides the results on the computation of bounds, this paper contains a characterization of the class of bounded nets having a unique possible steady-state behaviour, and a discussion of their ergodicity conditions. In particular, the concept of liveness bound for transitions is a new behavioural property, that comes directly from considerations related to the timing semantics of a timed Petri net model. It generalizes the usual concept of liveness for a transition, and provides an example of possible interleaving between qualitative and quantitative analysis for timed and stochastic Petri nets.

Extensions of the results presented in this paper are already being considered. In particular, work is in progress to study the case of live free choice nets, that can model queueing networks plus synchronization and concurrency. In this case, the idea is that several repetitive firing count vectors can be reproduced in steady-state, but the decisions are done only in places with more than one output transition, where the free choice net poses no restriction. Thus the selection is completely governed by the stochastic interpretation of the net, and an "average firing speed vector" can be defined independently of the marking of the net.

ACKNOWLEDGEMENTS

The authors are grateful to the anonymous referees for their detailed comments and suggestions on the draft of this paper.

References

[AMBB*89] 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 Transactions on Software Engineering*, July 1989.

[Bra83] G. W. Brams. *Réseaux de Petri: Théorie et Pratique. T.1. théorie et analyse.* Masson, Paris, 1983.

[BV84] E. Best and K. Voss. Free-choice systems have home states. *Acta Informatica*, (21):89–100, 1984.

[CCCS89] J. Campos, G. Chiola, J. M. Colom, and M. Silva. Tight polynomial bounds for steady-state performance of marked graphs. In *Proc. of 3rd Intern. Workshop PNPM'89*, IEEE-CS Press, Kyoto, Japan, December 1989.

[CHEP71] F. Commoner, A. W. Holt, S. Even, and A. Pnueli. Marked directed graphs. *Journal of Computer and System Science*, (5):511–523, 1971.

[Cox55] D. R. Cox. A use of complex probabilities in the theory of stochastic processes. *Proc. Cambridge Philosophical Society*, 51(2):313–319, April 1955.

[FN85] G. Florin and S. Natkin. Les réseaux de Petri stochastiques. Thesis de Doctorat d'Etat, Université P. et M. Curie, Paris, 1985.

[GP87] E. Gelenbe and G. Pujolle. *Introduction to Queuing Networks.* John Wiley, 1987.

[Har85] C. M. Harris. A note on mixed exponential approximations for GI/G/1 queues. *Computers and Operations Research*, (12), 1985.

[LR78] L. H. Landweber and E. L. Robertson. Properties of conflict-free and persistent Petri nets. *Journal of the ACM*, 25(3), April 1978.

[Mur77] T. Murata. Circuit theoretic analysis and synthesis of marked graphs. *IEEE Transactions on Circuits and Systems*, 24(7), July 1977.

[Mur83] K. G. Murty. *Linear Programming.* John Wiley & Sons, 1983.

[Mur85] T. Murata. Use of resource-time product concept to derive a performance measure of timed Petri nets. In *Proc. 1985 Midwest Symposium Circuits and Systems*, Louisville, U.S.A. August 1985.

[Ram74] C. Ramchandani. *Analysis of Asynchronous Concurrent Systems by Petri Nets.* PhD thesis, M.I.T. Cambridge, Mass. February 1974.

[RH80] C. V. Ramamoorthy and G. S. Ho. Performance evaluation of asynchronous concurrent systems using Petri nets. *IEEE Transactions on Software Engineering*, 6(5):440–449, September 1980.

[Sil87] M. Silva. Towards a synchrony theory for P/T nets. In K. Voss, H. Genrich, and G. Rozenberg, editors, *Concurrency and Nets*, pages 435–460, L.N.C.S. Springer-Verlag, Berlin, 1987.

[WT85] R. P. Wiley and R. R. Tenney. Performance evaluation of decision-free stochastic timed Petri nets. In *Proc. 24th Conf. Decision and Control*, Ft. Lauderdale, FL. December 1985.