Asymptotic Throughput of Continuous Timed Petri Nets

GUY COHEN[t‡]   STÉPHANE GAUBERT[‡*]   JEAN-PIERRE QUADRAT[‡*]

† Centre Automatique et Systèmes, École des Mines de Paris,
35 rue Saint-Honoré, 77305 Fontainebleau Cedex, France. e-mail: cohen@cas.ensmp.fr

‡ INRIA, Domaine de Voluceau, B.P. 105, 78153 Le Chesnay Cedex, France.
* e-mail: {Stephane.Gaubert,Jean-Pierre.Quadrat}@inria.fr

Abstract

We set up a connection between Continuous Timed Petri Nets (the fluid version of usual Timed Petri Nets) and Markov decision processes. We characterize the subclass of Continuous Timed Petri Nets corresponding to undiscounted average cost structure. This subclass satisfies conservation laws and shows a linear growth: one obtains as mere application of existing results for Dynamic Programming the existence of an asymptotic throughput. This rate can be computed using Howard-type algorithms, or by an extension of the well known cycle time formula for timed event graphs. We present an illustrating example and briefly sketch the relation with the discrete case.

Keywords— Petri Nets, Dynamic Programming, Markov Decision Processes, Discrete Event Systems, Max plus algebra.

I. Introduction

The fact that a subclass of Discrete Event Systems can be represented by linear equations in the (min,+) or in the (max,+) semiring is now almost classical [7], [2]. The (min,+) linearity allows the presence of synchronization and saturation features but prohibits the modeling of many interesting phenomena such as "birth" and "death" processes (multiplication of tokens) and concurrency. More recently [3], [8], it was realized that under some assumptions on the routing policies, these additional features could be represented by more general recurrences, involving both conventional linear systems and (min,+) linear systems. From the control theoretical point of view, these are *polynomial* systems over the (min,+) algebra, that is, the exact (min,+) counterpart of conventional polynomial discrete time systems. This approach was outlined in [8], where in particular the (min,+) analogue of Volterra expansion was given.

Another (simpler) point of view is based on Markov decision processes. As shown in [8], the "polynomial" Petri Net equations can be interpreted as the dynamic programming equations of a canonical Markov decision process associated with the net, equipped with an additive discounted cost. More explicitely, the counter function (number of firings) of transition q at time t is equal to the value function at state q for the associated Markov decision process in horizon t. Of particular interest is the case of undiscounted costs: then the value function grows linearly as a function of the horizon, and for the corresponding Petri Nets (that we call undiscounted), there exists an *asymptotic throughput* (mean number of firings of a given transition per time unit). Undiscounted Petri Nets are characterized by the following simple structural property: there are as many input as output arcs at each place. Then, the routing policy is uniquely defined (one has to route the tokens equally towards downstream arcs). We show that undiscounted Petri Nets admit P-invariants (linear combination of markings invariant by firing of transitions). They also admit T-invariants (sequences of firings preserving the marking) which are best interpreted as an input-output *homogeneity* property: if we distinguish between input transitions (representing e.g. the availability of raw materials) and output transitions (representing finished parts) and if we add one unit of each input material, then one obtains one more unit of each finished part. Finally, we introduce the class of Petri Nets *with potential*, obtained from undiscounted Petri Nets via rescalings (changes of units). This class is suited to the modeling of production systems in which parts are produced according to different ratios.

A more complete presentation in a system theoretical spirit will be found in [8], to which the reader is referred for omitted proofs. Here, we focus on the main computational consequence of this approach, that is, the asymptotic throughput formula.

## II. Continuous Timed Petri Nets under Stationary Routing Policies

Fig. 1. A Timed Petri Net.

The definition of Continuous Timed Petri Nets is syntaxically very similar to that of conventional Timed Petri Nets: one has to specify the topology of the graph, the initial marking and the durations. The main difference lies in

the functioning and the interpretation of the system, since fluids instead of tokens circulate in the net.

Definition II.1 (CTPN). A Continuous Timed Petri Net with Multipliers (CTPN) is a valued bipartite graph given by a 5-tuple $\mathcal{N} = (\mathcal{P}, \mathcal{Q}, M, m, \tau)$ with the following characteristics.

1. $\mathcal{P}$ is a finite set whose elements are called *places*. Places should be seen as reservoirs, with input and output pipes, in which a liquid flows according to a dynamics described later on.

2. $\mathcal{Q}$ is a finite set whose elements are called *transitions*. Transitions mix the flows coming from the places immediately upstream in given proportions and instantaneously[1], and pour the resulting liquid in downstream places also in given proportions.

3. $M \in (\mathbb{R}^+)^{\mathcal{P} \times \mathcal{Q} \cup \mathcal{Q} \times \mathcal{P}}$. The *multiplier* M_{pq} (resp. M_{qp}) gives the number of edges from transition q to place p (resp. from place p to transition q). We allow non integer number of edges. The zero value for M codes the absence of edge. We say that vertex (place or transition) r is *upstream* vertex s if $M_{sr} \neq 0$. Equivalently, s is *downstream* r. We denote by r^{out} the set of vertices downstream vertex r and by r^{in} the set of upstream vertices. Multipliers determine the mixing and dispatching proportions as follows: transition q takes M_{qp} molecules[2] of fluid from each upstream place p, and produces $M_{p'q}$ molecules[3] of fluid in each downstream place p'. The mixing process at transition q continues as long as *all* the upstream places are non empty. When a place is upstream *several* transitions, we assume that there is a routing mechanism fixing which proportions of the flow should be sent to the concurrent downstream transitions. Keeping the discrete terminology, we will still call *firing* of transition q the consumption of M_{qp} molecules in each upstream place p and the production of $M_{p'q}$ molecules in each downstream place, but now, transition firings are counted with real numbers.

4. $m \in (\mathbb{R}^+)^{\mathcal{P}}$ represents the initial marking: m_p gives the amount of fluid initially available in place p.

5. $\tau \in (\mathbb{R}^+)^{\mathcal{P}}$ (holding times): τ_p gives the sojourn time in place p, i.e. the minimal time from the entry of a molecule in place p to its avaibility for the firing of downstream transitions. This delay may be caused for example by a preparation time required for heating or homogenizing the fluid.

---

[1] There is no loss of modelling power in assuming the mixing operation to be instantaneous, since mixing delays can be incorporated in sojourn times in reservoirs, possibly after adding places and transitions.

[2] Continuous Petri nets differ from discrete (conventional) ones in that the fluid quantities are infinitely divisible. E.g. consider a transition q with two upstream places p_1, p_2, one downstream place p_3, and mixing ratios $M_{qp_1} = 1, M_{qp_2} = 2, M_{p_3q} = 1$. Assume that there is 1 molecule of fluid in each upstream place. Then the continuous approximation consumes $1/2$ molecule in p_1 and 1 molecule in p_2 to produce $1/2$ molecule in p_3, while in the discrete —more realistic— case, the transition is blocked (no chemical reaction can occur).

[3] Quantities of fluid are measured in number of molecules rather than in volume or mass. Note that in general, fluid measures are not preserved by mixing operations, i.e. $\sum_{q \in p^{\text{in}}} M_{pq} \neq \sum_{q' \in p^{\text{out}}} M_{q'p}$.

Definition II.2 (Routing Policy). A (stationary, origin independent[4]) *routing policy* is a map $\rho : \mathcal{Q} \times \mathcal{P} \to \mathbb{R}^+$, such that $\forall p \in \mathcal{P}$, $\sum_{q \in p^{\text{out}}} \rho_{qp} = 1$. ρ_{qp} determines the proportion of fluid routed to the downstream transition q by place p. The initial stock of fluid m_p is routed with the same proportions.

We next give the dynamic equations satisfied by the Timed Petri Net. *Counter functions* are associated with nodes of the graph:

1. $Z_p(t)$ denotes the cumulated quantity of fluid[5] which has entered place p up to time t, including the initial stock;
2. $Z_q(t)$ denotes the cumulated number of firings of transition q up to time t.

All these *cumulated* quantities are of course nondecreasing functions.

We focus here on the *autonomous* regime, that is we assume that the counter trajectories $Z_r(t)$, $r, s \in \mathcal{Q} \cup \mathcal{P}$ are frozen for $t < 0$ at a given initial condition, and we let the system evolve freely for $t \geq 0$. A more general input-output approach (the set of transitions being partitionned in input/state/output transitions) is detailed in [8], [9].

We introduce the notation

$$\mu_{pq} \stackrel{\text{def}}{=} M_{pq}, \quad \mu_{qp} \stackrel{\text{def}}{=} M_{qp}^{-1}, \quad \mu'_{qp} \stackrel{\text{def}}{=} \mu_{qp} \rho_{qp} .$$

Proposition II.3. *The counter variables of a CTPN satisfy the following equations[6]*

$$Z_q(t) = \min_{p \in q^{\text{in}}} \mu'_{qp} Z_p(t - \tau_p) , \qquad (1a)$$

$$Z_p(t) = m_p + \sum_{q \in p^{\text{in}}} \mu_{pq} Z_q(t) . \qquad (1b)$$

Remark II.4. From (1), we deduce the transition-to-transition equation:

$$Z_q(t) = \min_{p \in q^{\text{in}}} \left[\mu'_{qp} \left(m_p + \sum_{q' \in p^{\text{in}}} \mu_{pq'} Z_{q'}(t - \tau_p) \right) \right] . \quad (2)$$

Dually, the place-to-place equation can be obtained:

$$Z_p(t) = m_p + \sum_{q' \in p^{\text{in}}} \mu_{pq'} \min_{p' \in (q')^{\text{in}}} \left(\mu'_{q'p'} Z_{p'}(t - \tau_{p'}) \right) .$$

$$(3)$$

Remark II.5. If $\tau_p = 0$ for some places, system (1) becomes implicit and we may have difficulties in proving the existence of a finite solution. We say that the CTPN is *explicit* if there is no circuits containing only places with zero holding times. Then, (1) becomes explicit in time, that is, if one knows the past values $Z_r(\tau), \tau < t$, the counter functions $Z_s(t), s \in \mathcal{P} \cup \mathcal{Q}$, can be computed sequentially from (1) (using an appropriate order of the vertices).

---

[4] More general routing policies were considered in [8]. It is possible, up to a minor increase of complexity, to consider *origin dependent* routings, that is to assume that the routing of a given quantity of fluid at place p depends of the transition from which it comes. Initial stocks may also admit a particular routing (distinct from the stationary one). We shall not enter in these details here.

[5] Measured in number of molecules.

[6] We adopt the convention $\sum_{q \in \emptyset}() = 0$, so that (1b) becomes $Z_p(t) = m_p$ when $p^{\text{in}} = \emptyset$.

Remark II.6. We note that (1) reads as the coupling of a conventional linear system with a $(\min, \times)$ linear system. Let us assume $\tau_p = 1, \forall p$, for simplicity. Then[7],

$$Z_Q(t) = \mu'_{QP} \otimes Z_P(t-1) , \quad (4a)$$

$$Z_P(t) = m + \mu_{PQ} Z_Q(t) , \quad (4b)$$

where $(A \otimes x)_i = \bigoplus_j A_{ij} \otimes x_j = \min_j A_{ij} x_j$ is the matrix product of the dioid[8] $\mathbb{R}_{\min,\times} \stackrel{\text{def}}{=} (\mathbb{R}^{+*} \cup \{+\infty\}, \min, \times)$.

Example II.7. The Timed Petri Net depicted in Fig. 1 has four places $\mathcal{P} = \{p_1, \cdots, p_4\}$ and six transitions $\mathcal{Q} = \{q_1, \cdots, q_6\}$. The coefficients M_{pq} and M_{qp} are visualized on the picture by the number of arcs going from a transition to a place or from a place to a transition. The initial stocks are not indicated. In order to simplify the notation, we shall write ρ_{ij} instead of $\rho_{q_i p_j}$, m_i instead of m_{p_i}, etc. The routing policy is displayed in the figure with greek letters, i.e. $\rho_{11} = \gamma$, $\rho_{41} = \delta$, $\rho_{32} = \alpha$, $\rho_{22} = \beta$, $\rho_{53} = 1$, $\rho_{14} = \eta$, $\rho_{64} = \zeta$, with $\alpha + \beta = 1$, $\gamma + \delta = 1$, $\eta + \zeta = 1$. Therefore, the dynamics of the system (see (4)) is described by the four matrices $m = (m_i)_{i=1\ldots4}, \tau = (\tau_i)_{i=1\ldots4}$,

$$\mu_{PQ} = \begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 2 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix},$$

$$\mu'_{QP} = \begin{pmatrix} \gamma & \infty & \infty & \eta \\ \infty & \beta/2 & \infty & \infty \\ \infty & \alpha & \infty & \infty \\ \delta & \infty & \infty & \infty \\ \infty & \infty & 1 & \infty \\ \infty & \infty & \infty & \zeta \end{pmatrix}.$$

The behavior of the system can be easily simulated by computing the matrix products (4), starting from an initial condition $Z_Q(t), t < 0$.

We refer the reader to [3] for the study of similar equations when the delays are random variables.

III. Stochastic Control Interpretation of Continuous Timed Petri Nets

We associate with a CTPN the following canonical Markov Decision Process.

1. The time evolves backward.
2. The set of Markov chain states is $\mathcal{Q}$. The probability $P^p_{qq'}$ of the transition $q \to q'$ from time n to time $n-1$ under the decision p is given by

$$P^p_{qq'} = \alpha^{-1}_{qp} \mu'_{qp} \mu_{pq'} \text{ with } \alpha_{qp} \stackrel{\text{def}}{=} \sum_{q' \in p^{\text{in}}} \mu'_{qp} \mu_{pq'} .$$

It is essential to note that the process moves *backward* in the graph, i.e. if the Petri Net has an arc $q' \to p \to q$, then the move $q \to q'$ for the associated Markov chain is allowed.

3. The set $\mathcal{P}_{\text{ad}}$ of admissible control histories is the set of sequences $p_1, \ldots, p_t$ such that $p_n \in q^{\text{in}}_n$ and the decision p_n is a feedback over q_n.
4. The discounted cost to be minimized knowing that we start at time t from state q is

$$J(p, t, q) = \mathbb{E}\{Z_{q_0}(0)\beta_{0t} + \sum_{n=1}^{t} \nu_{q_n p_n} \beta_{nt} | q_t = q\} ,$$

where we have used $\nu_{qp} \stackrel{\text{def}}{=} \mu'_{qp} m_p$ and denoted the state trajectory and control dependent actualization by

$$\beta_{st} \stackrel{\text{def}}{=} \prod_{j=s+1}^{t} \alpha_{q_j p_j} .$$

Proposition III.1. *For a CTPN such that* $\tau_p \equiv 1$, *the counter function coincides with the value function:*

$$Z_q(t) = \inf_{p \in \mathcal{P}_{\text{ad}}} J(p, t, q) . \quad (5)$$

Proof. If we write down the dynamic programming recursion for $V_q(t) \stackrel{\text{def}}{=} \inf_{p \in \mathcal{P}_{\text{ad}}} J(p, t, q)$, we obtain precisely Eq. (2). $\square$

The case of integer but non identically 1 delays could be interpreted along the same lines, although writing precisely the resulting Markov Decision Process would be slightly more involved.

Definition III.2. We say that a CTPN

1. is *balanced* if $\forall p, \sum_{q \in p^{\text{out}}} M_{qp} = \sum_{q' \in p^{\text{in}}} M_{pq'}$, that is if there are as many arcs upstream and downstream each place;
2. is *undiscounted* if $\alpha_{qp} \equiv 1$;
3. admits a *potential* if there exists a vector $v \in (\mathbb{R}^{+*})^{\mathcal{Q}}$ (called potential) such that the change of variable $Z_q = v_q Z'_q$ makes the CTPN undiscounted.

Theorem III.3. *1. A CTPN becomes undiscounted under a (stationary, origin independent) routing policy iff it is balanced.*

2. A CTPN admits a potential v for some (stationary, origin independent) routing iff

$$\forall p, \sum_{q \in p^{\text{out}}} v_q M_{qp} = \sum_{q' \in p^{\text{in}}} M_{pq'} v_{q'} . \quad (6)$$

Then,

$$\forall q \in Q, \ p \in q^{\text{in}}, \ v_q = \sum_{q' \in p^{\text{in}}} \mu'_{qp} \mu_{pq'} v_{q'} . \quad (7)$$

Proof. Since the specialization of item 2 to unit potential yields item 1, we only prove item 2. The CTPN has potential v iff for all p, the matrix

$$P^p_{qq'} = v^{-1}_q M^{-1}_{qp} \rho_{qp} M_{pq'} v_{q'}$$

is stochastic. Summing up as $q' \in p^{\text{in}}$, we get

$$v_q M_{qp} = \sum_{q' \in p^{\text{in}}} \rho_{qp} M_{pq'} v_{q'} , \quad (8)$$

that is (7). Summing up (8) as $q \in P^{\text{out}}$ and using $\sum_{q \in P^{\text{out}}} \rho_{qp} = 1$ we get the necessary condition (6). Conversely, when (6) is true, the routing policy

$$\rho_{qp} = \frac{v_q M_{qp}}{\sum_{q' \in P^{\text{out}}} v_{q'} M_{q'p}} \qquad (9)$$

turns $v_q^{-1} M_{qp}^{-1} \rho_{qp} M_{pq'} v_{q'}$ to a stochastic matrix which shows that (6) is also a sufficient condition. $\qquad \square$

Example III.4. The Petri Net shown in Fig 1 is balanced since at each place the same number of arcs arrives and leaves. Therefore the Petri Net admits an undiscounted stochastic control interpretation as soon as the routing policy assigns the flow equally to the outgoing arcs, that is

$$\alpha = 1/3, \ \beta = 2/3, \ \gamma = \delta = 1/2, \ \eta = \zeta = 1/2 . \quad (10)$$

Indeed, in this case the potential is $v_q \equiv 1$, and the only compatible routing is given by (9). The non-zero rows of the matrices P^p are given as follows:

$$
\begin{array}{c}
P_1^1 \\
P_1^4 \\
P_2^2 \\
P_3^2 \\
P_4^1 \\
P_5^3 \\
P_6^4
\end{array}
\begin{pmatrix}
0 & 1/2 & 0 & 1/2 & 0 & 0 \\
0 & 0 & 0 & 0 & 1/2 & 1/2 \\
1/3 & 0 & 2/3 & 0 & 0 & 0 \\
1/3 & 0 & 2/3 & 0 & 0 & 0 \\
0 & 1/2 & 0 & 1/2 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1/2 & 1/2
\end{pmatrix}
\qquad (11)
$$

Example III.5. Using $\alpha + \beta = 1, \gamma + \delta = 1, \eta + \zeta = 1$ together with (7), some elementary elimination shows that the CTPN admits a potential as soon as $\alpha = 1/3$. Then:

$$v = \begin{pmatrix} 1 & 1 & 1 & (1-\gamma)/\gamma & 1 & (1-\eta)/\eta \end{pmatrix} . \quad (12)$$

IV. Asymptotic Properties of Undiscounted Petri Nets

In this section we use the stochastic control interpretation to obtain explicit formulas for the throughput of balanced CTPN and CTPN with potential.

Theorem IV.1. *For a strongly connected undiscounted CTPN, we have*

$$\lim_{t \to \infty} \frac{1}{t} Z_q(t) = \lambda, \quad \forall q ,$$

where λ is a constant. The periodic throughput λ is characterized as the unique value for which a finite vector w is solution of

$$w = \min_p \left(\nu_{\cdot p} - \lambda \tau_p + P^p w \right) . \quad (13)$$

Proof. This is an adaptation of standard stochastic control results [12, Chap. 33, Th. 4.1]. The growth rate λ is independent of the initial point q for the subclass of communicating systems[9]. This assumption is equivalent to the strong connectedness of the net. $\qquad \square$

[9]The system is communicating if for all q, q', there is a policy u and an integer k such that $(P_{qq'}^u)^k > 0$, that is, there is a policy such that there exists a path from q to q' with non zero probability.

A *feedback policy* (or policy[10], for short) is a map $u : \mathcal{Q} \to \mathcal{P}$. The policy is *admissible* if $u(q) \in q^{\text{in}}$, that is, if setting $p_n = u(q_n)$ yields an admissible policy for the corresponding stochastic control problem. The following vectors and matrices are associated with a policy u.

$$\nu_q^u \overset{\text{def}}{=} \nu_{qu(q)} , \quad \tau_q^u \overset{\text{def}}{=} \tau_{u(q)} , \quad P_{qq'}^u \overset{\text{def}}{=} P_{qq'}^{u(q)} .$$

Remark IV.2. Equation (13) can be solved efficiently using Howard algorithm (policy iteration).

1. Initialization: select a policy u such that P^u is strongly connected (irreducible).
2. Given a policy u^n solve

$$w^n = \nu^{u^n} - \lambda^n \tau^{u^n} + P^{u^n} w^n ,$$

where the unknowns are (λ^n, w^n). This is a linear system which has a unique solution if we impose $w_0^n = 0$ (for example) and if the policy matrix P^{u^n} is strongly connected.
3. Given (λ^n, w^n) improve the policy by

$$u^{n+1}(q) = \arg\min_p [\nu_{\cdot p} - \lambda^n \tau_p + P^p w^n]_q , \forall q \in \mathcal{Q} .$$

This algorithm defines a positive decreasing sequence of λ^n converging in a finite number of steps to the searched throughput λ. The algorithm is properly defined if at each step we get a strongly connected policy matrix. This is not always possible (in particular, for the example dealt with here, the only two policy matrices are not strongly connected, see Ex.IV.4 infra). We will not discuss here the extension of the algorithm to the non strongly connected case.

There is an equivalent characterization of λ which exhibits the analogy with the Timed Event Graphs case in a better way. We denote by $\mathcal{R}(u)$ the set of final classes of the matrix P^u. For each class $r \in \mathcal{R}(u)$, we have a unique invariant measure π^{ur} with support r (i.e. $\pi^{ur} P^u = \pi^{ur}$, and $\pi_q^{ur} = 0$ if $q \notin r$.)

Theorem IV.3 (Asymptotic Throughput Formula).
For an undiscounted CTPN, we have

$$\lambda = \min_u \min_{r \in \mathcal{R}(u)} \frac{\pi^{ur} \nu^u}{\pi^{ur} \tau^u} . \qquad (14)$$

Thus, λ is the minimal ratio of the mean marking over the mean holding time in the places visited while following a stationary policy. This result should be compared with the well know periodicity theorem (see e.g. [2]) which states that a (strongly connected) Timed Event Graph reaches a periodic regime after a finite time, with throughput

$$\lambda = \min_c \frac{\sum_{p \in c} m_p}{\sum_{p \in c} \tau_p} , \qquad (15)$$

where the minimum is taken over the (elementary) circuits of the graph. Since in the case of Timed Event Graphs, the final classes are precisely circuits and the invariant measures are uniform on the final classes, (14) reduces to (15).

[10]This feedback policy has nothing to do with the *routing* policy introduced in section II.

2032

Example IV.4. For the current example, it is easy to compute the throughput introduced in Th. IV.1, which is interpreted in stochastic control terms as:

$$\lambda = \min_{\substack{t=0,\cdots,T-1 \\ p_t \in \{p_1, p_4\}}} \lim_{T \to \infty} \frac{1}{T} \mathbb{E} \sum_{0, T-1} \nu_{q_t p_t}.$$

Indeed, they are only two policies. If we use the strategy u_1: "choose p_1 when we are in q_1" (see Fig 2), we obtain the matrix

$$P^{u_1} = \begin{pmatrix} 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 1/3 & 0 & 2/3 & 0 & 0 & 0 \\ 1/3 & 0 & 2/3 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 \end{pmatrix}.$$

If we use the strategy u_4: "choose p_4 when we are in q_1"

Fig. 2. Markov chain associated with P^{u_1}.

(see Fig 3),, we get the matrix,

$$P^{u_4} = \begin{pmatrix} 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 1/3 & 0 & 2/3 & 0 & 0 & 0 \\ 1/3 & 0 & 2/3 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 \end{pmatrix}.$$

The instantaneous cost at each step is given by the first row

Fig. 3. Markov chain associated with P^{u_4}.

of the following matrix if the decision p_1 is taken, and by the second one if the decision p_4 is taken:

$$\nu' = \begin{pmatrix} m_1/2 & m_2/3 & m_2/3 & m_1/2 & m_3 & m_4/2 \\ m_4/2 & m_2/3 & m_2/3 & m_1/2 & m_3 & m_4/2 \end{pmatrix}.$$

The only final class of the Markov chain for the first strategy is $\{q_1, q_2, q_3, q_4\}$ and we get the invariant measure of probability

$$\pi^1 = (\ 1/5 \quad 1/5 \quad 2/5 \quad 1/5 \quad 0 \quad 0\ ).$$

For the second strategy, the only final class of the Markov chain is $\{q_1, q_5, q_6\}$ and we get the invariant measure of probability

$$\pi^4 = (\ 1/3 \quad 0 \quad 0 \quad 0 \quad 1/3 \quad 1/3\ ).$$

Therefore the throughput of the system is given by:

$$\lambda = \min \left(\frac{m_1 + m_2}{2\tau_1 + 3\tau_2}, \frac{m_3 + m_4}{\tau_3 + 2\tau_4} \right). \quad (16)$$

A simulation of the system for $m = (\ 0 \quad 3 \quad 10 \quad 10\ )$ and $\tau = (\ 1 \quad 1 \quad 2 \quad 2\ )$ is shown in Fig 4. We note that the minimum in (16) is equal to 0.6 and that it is attained for the first term corresponding to the "critical" final class $\{q_1, q_2, q_3, q_4\}$. The transient regime for Z_5, Z_6 can be explained easily by the fact that the non critical transitions q_5, q_6 first consume the initial stock $m_3 = m_4 = 10$ at their own regime, before being delayed by the critical final class.


Fig. 4. Continuous behavior

It is not surprising that the terms appearing in (14) are indeed *invariants* of the net.

Theorem IV.5 (Invariants). *Given an undiscounted CTPN, for all policy u and for all final class r associated with u,*

$$I^{ur} \stackrel{\text{def}}{=} \pi^{ur} \nu^u = \sum_{q \in r} \pi^{ur} \nu^u_q \quad (17)$$

is invariant by firing transitions.

Theorem IV.6 (Homogeneity). *If $Z(t), t \in \mathbb{R}$, is an admissible trajectory of an undiscounted CTPN, then[11] $Z(t) + 1, t \in \mathbb{R}$, is also an admissible trajectory.*

This theorem should be interpreted as follows: if some transitions can be regarded as input ports of raw material (with no upstream arcs), then an increase of 1 unit of all inputs induce the same increase of all the other quantities.

These results can be extended to the case of CTPN with potential. With a feedback policy u we associate the matrix R^u: $R^u_{qq'} = \mu'_{qu(q)} \mu_{u(q)q'}$ if $q' \in u(q)^{\text{in}}$ ($R^u_{qq'} = 0$ otherwise); we denote by $\mathcal{R}(u)$ the set of final classes of R^u; which each final class r we associate a left eigenvector of R^u: $\pi^{ur} = \pi^{ur} R^u$ with support r; and we define ν^u, τ^u as in Theorem IV.3. We denote by diag v the diagonal

---

[11] We denote by $Z(t) + 1$ the vector $(Z_r(t) + 1)_{r \in \mathcal{P} \cup \mathcal{Q}}$

matrix with diagonal entries $(\mathrm{diag}\,v)_{qq} = v_q$. Then, the following formula is an immediate consequence of Theorem IV.3.

Corollary IV.7. *For a strongly connected CTPN with potential v, we have* $\lim_{t\to\infty} \frac{1}{t} Z_q(t) = \lambda_q$ *, where*

$$v_q^{-1}\lambda_q = \min_u \min_{r\in\mathcal{R}(u)} \frac{\pi^{ur}\nu^u}{\pi^{ur}(\mathrm{diag}\,v)\tau^u} \,. \tag{18}$$

Example IV.8. When $\alpha = 1/3$, the Petri Net of Fig 1 admits the potential (12). An immediate application of (18) yields:

$$\lambda = \min\left(\frac{(m_1+m_2)\gamma}{\tau_1+3\gamma\tau_2}, \frac{(m_3+m_4)\eta}{\tau_4+\eta\tau_3}\right)\,.$$

V. Discrete vs. Continuous Behavior

To conclude, we would like to indicate how these results may help the analysis of conventional (discrete) Timed Petri Nets.

Firstly, let us recall how the dynamics (1) has to be modified in the discrete case. We now call *routing policy* at place p a family $\{\Pi_{qp}\}_{q\in p^{out}}$, where Π_{qp} is a nondecreasing map $\mathbb{N} \to \mathbb{N}$: $\Pi_{qp}(n)$ *tells the number of tokens routed to q from p among the first n ones.* Since $\{\Pi_{qp}\}_{q\in p^{out}}$ is a partition of the flow from p, we have: $\forall n$, $\sum_{q\in p^{out}}\Pi_{qp}(n) = n$. It is not difficult to see that the transition-to-transition equation (2) becomes

$$Z_q(t) = \min_{p\in q^{in}}\left\lfloor \mu_{qp}\left(\sum_{q'\in p^{in}}\Pi_{qp}\left(m_p + \mu_{pq'}Z_{q'}(t-\tau_p)\right)\right)\right\rfloor\,, \tag{19}$$

where $\lfloor x\rfloor = \sup\{n\in\mathbb{N} \mid n\le x\}$. See [8] for more details.

The discrete counterpart of a stationary routing with (rational) ratio ρ_{qp} is obtained for a periodic function Π_{qp} with slope ρ_{qp}. i.e. we assume that there is an integer c (periodicity) such that $\Pi_{qp}(n + c) = \Pi_{qp}(n) + \rho_{qp}c$.

Example V.1. A possible discrete version of the routing used in Ex. 10 at place p_2 is the following:

$$\Pi_{32}(n) = \left\lfloor\frac{n}{3}\right\rfloor\,,\quad \Pi_{22}(n) = \left\lfloor\frac{n+1}{3}\right\rfloor + \left\lfloor\frac{n+2}{3}\right\rfloor\,.$$

In simpler terms, the tokens numbered $3k$ are routed to transition q_3, while the tokens numbered $3k+1$ or $3k+2$ are routed to transition q_2. Other choices of congruence are of course possible. Similarly, we take:

$$\Pi_{11}(n) = \Pi_{14}(n) = \left\lfloor\frac{n}{2}\right\rfloor\,,$$

$$\Pi_{41}(n) = \Pi_{64}(n) = \left\lfloor\frac{n+1}{2}\right\rfloor\,.$$

A simulation for the same values as in Ex IV.4 is shown in Fig 5. The asymptotic throughput (obtained experimentally) is $\lambda = 0.5$, to be compared with the continuous throughput 0.6.


Fig. 5. Discrete behavior

More generally, the following results can be stated. Firstly, we can obviously bound $\Pi_{qp}(n)$ by affine functions, i.e. $\rho_{qp}\times(n+\underline{m}_p) \le \Pi_{qp}(n) \le \rho_{qp}\times(n+\overline{m}_p)$ for some $\underline{m}_p, \overline{m}_p$. Then, a comparison of (2) and (19) shows that the discrete system is bounded from above by the associated CTPN with marking $m'_p = m_p + \overline{m}_p$, and from below by the associated CTPN with marking $m''_p = m_p + \underline{m}_p - 1$. Secondly, consider a strongly connected balanced TPN such that ρ_{qp} is uniform (i.e. it routes tokens to all the downstream arcs in the same proportions), and assume the existence of a periodic throughput $\lambda_q \overset{def}{=} \lim_{t\to\infty} t^{-1}Z_q(t)$ for all t. Then, some reworking of the proof of Theorem IV.1 shows that all the transitions admit the same rate: $\lambda_q = \lambda_{q'}, \forall q, q'$.

References

[1] F. Baccelli, G. Cohen, and B. Gaujal. Recursive equations and basic properties of timed Petri nets. *J. of Discrete Event Dynamic Systems*, 1(4):415–439, 1992.

[2] F. Baccelli, G. Cohen, G.J. Olsder, and J.P. Quadrat. *Synchronization and Linearity*. Wiley, 1992.

[3] F. Baccelli, S. Foss, and B. Gaujal. Structural, temporal and stochastic properties of unbounded free-choice Petri nets. Rapport de Recherche 2411, INRIA, 1994.

[4] J. Le Bail, H. Alla, and R. David. Hybrid Petri nets. In *Proc. of the ECC 91*, Grenoble, France, July 1991.

[5] P. Chretienne. *Les Réseaux de Petri Temporisés*. Thèse Université Pierre et Marie Curie (Paris VI), Paris, 1983.

[6] G. Cohen, D. Dubois, J.P. Quadrat, and M. Viot. Analyse du comportement périodique des systèmes de production par la théorie des dioïdes. Rapport de recherche 191, INRIA, Le Chesnay, France, 1983.

[7] G. Cohen, P. Moller, J.P. Quadrat, and M. Viot. Algebraic tools for the performance evaluation of discrete event systems. *IEEE Proceedings: Special issue on Discrete Event Systems*, 77(1), Jan. 1989.

[8] G. Cohen, S. Gaubert, and J.P. Quadrat. Algebraic System Analysis of Timed Petri Nets. In *Idempotency*, J. Gunawardena Ed. Collection of the Isaac Newton Institute, Cambridge University Press, to appear.

[9] G. Cohen, S. Gaubert, and J.P. Quadrat. Timed Event Graphs with Multipliers and Homogeneous Min-Plus Systems. *submitted for publication*.

[10] M. Plus. A linear system theory for systems subject to synchronization and saturation constraints. In *Proceedings of the first European Control Conference*, Grenoble, July 1991.

[11] P.J. Schweitzer and A. Federgruen. Geometric convergence of value-iteration in multichain Markov decision problems. *Adv. Appl. Prob.*, 11:188–217, 1979.

[12] P. Whittle. *Optimization over Time*, volume II. Wiley, 1986.

[13] W.H.M. Zijm. Generalized eigenvectors and sets of nonnegative matrices. *Lin. Alg. and Appl.*, 59:91–113, 1984.