



HAL
open science

Throughputs in stochastic free-choice nets, existence, computations and optimizations

Anne Bouillard, Bruno Gaujal, Jean Mairesse

► **To cite this version:**

Anne Bouillard, Bruno Gaujal, Jean Mairesse. Throughputs in stochastic free-choice nets, existence, computations and optimizations. 44-th IEEE Conference on Decision and Control, 2005, Sevilla, Spain. pp.1-19. inria-00071380

HAL Id: inria-00071380

<https://inria.hal.science/inria-00071380v1>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Throughputs in stochastic free-choice nets,
existence, computations and optimizations***

Anne Bouillard — Bruno Gaujal — Jean Mairesse

N° 5888

Mars 2006

Thème NUM



***Rapport
de recherche***



Throughputs in stochastic free-choice nets, existence, computations and optimizations

Anne Bouillard^{*}, Bruno Gaujal[†], Jean Mairesse[‡]

Thème NUM — Systèmes numériques
Projet MESCAL

Rapport de recherche n° 5888 — Mars 2006 — 19 pages

Abstract: In this paper, live and bounded free-choice Petri nets with stochastic firing times are considered. Several classical routing policies, namely the race policy, Bernoulli routings, and periodic routings, are compared in terms of the throughputs of the transitions. First, under general i.i.d. assumptions on the firing times, the existence of the throughput for the three policies is established. We also show that the ratio between the throughputs of two transitions depend only on the asymptotic frequencies of the routings, and not on the routing policy. On the other hand, the total throughput depends on the policy, and is higher for the race policy than for Bernoulli routings. Second, we show how to compute the throughput for exponentially distributed free-choice nets under the three policies. This is done by using Markov processes over appropriate state spaces. We use this to compare the performance of periodic and Bernoulli routings. Finally, we derive optimal policies under several information structures, namely, the optimal pre-allocation, the optimal allocation, and the optimal non-anticipative policy.

Key-words: Free Choice Petri Nets, Race Policy, Optimal Routing Policy.

^{*} LIP, ENS-Lyon, 46 Allée d'Italie, 69364 Lyon cedex07, France (Anne.Bouillard@ens-lyon.fr)

[†] INRIA, Laboratoire ID-IMAG, Mescal Project INRIA-UJF-CNRS-INPG, 51, avenue Jean Kuntzmann, F-38330 Montbonnot, France (Bruno.Gaujal@imag.fr)

[‡] Université Paris 7, LIAPA - Case 7014, 2 place Jussieu, 75251 Paris Cedex 05, France

Débit dans les réseaux à choix libres stochastiques : existence, calculs et optimisations.

Résumé : Dans cet article, nous considérons des réseaux de Petri vivants bornés avec des temps de tir stochastiques. Nous comparons plusieurs politiques de routage classiques : la politique de compétition, les routages Bernoulli et les routages périodiques. Pour ce faire nous étudions le débit des transitions du réseau. Tout d'abord, pour des temps de tir avec des distributions i.i.d. générales, nous établissons l'existence du débit pour ces trois politiques. Nous montrons aussi que le rapport entre les débits de deux transitions dépend seulement des fréquences de routage et non de la politique de routage. En revanche, le débit total dépend de la politique de résolution de conflits choisie : il est plus élevé pour la politique de compétition que pour le routage Bernoulli. Dans un deuxième temps, nous montrons comment calculer ce débit pour des temporisations exponentiellement distribuées pour ces trois politiques de routage en utilisant des processus de Markov sur des espaces d'états appropriés. Nous utilisons ces méthodes pour comparer les routages Bernoulli et périodiques. Enfin, nous montrons comment trouver les politiques qui maximisent le débit en fonction de l'information disponible. En utilisant des processus de décision de Markov nous exhibons les politiques de pré-allocation, d'allocation et non-anticipative optimales.

Mots-clés : Réseaux de Petri à choix libres, Compétition, Stratégie de Routage Optimal.

1 Introduction

In this paper, we consider a live and bounded free-choice net with stochastic firing times and we analyze classical policies of conflict resolution in terms of the throughput of the transitions (number of firings per second). The first policy is the famous *race policy*, see for instance [1]. The other policies are Bernoulli routings, periodic routings, and throughput-optimal routings.

This problem has already been considered for timed deterministic *fluid* Petri nets. Two different models of fluid Petri nets have been studied, in [9] and [15]. In both cases, it has been proved that the throughput is simply the solution of a linear program ([10, 15]). The discrete case is more involved. The *deterministic discrete* free-choice case has been studied in [6] and has a high combinatorial complexity. On the other hand, the existence of the throughput for *stochastic* free-choice nets with general i.i.d. firing times and Bernoulli routings is established in [11] but no means of computation is provided.

Here, we first show the existence of the throughput for the race policy and the periodic routings for general i.i.d. timings. Then we compare the throughput obtained under the different policies, for a fixed asymptotic frequency of the routings. Let λ^k , $k \in \{\text{race}, \text{Ber}, \text{per}\}$, be the vector of the throughputs at the different transitions. We prove that there exists a vector v only depending on the asymptotic routing frequencies and such that $\lambda^k = \alpha^k v$, for $\alpha_k \in \mathbb{R}_+$.

In the second part of the paper, we show how to compute explicitly the throughput for exponentially distributed free-choice nets with Bernoulli routings, periodic routings and for the race policy. The race policy case is standard: the marking evolves as a continuous-time jump Markov process. As for Bernoulli and periodic routings, we construct a Markov process which is not evolving on the marking reachability graph but on an extended state space which takes into account the possible routings. We show how to choose the parameters of the Bernoulli routing in order to maximize the throughput. We use these computations to compare Bernoulli routings with periodic routings. Numerical evidence suggests that balanced periodic routings provide better throughputs than Bernoulli routings, much like in open systems [2] or closed deterministic ones [8].

In the final part of the paper, we consider optimal policies. Observe that the race policy can be seen as a greedy policy which is locally optimal. Using Markov Decision Processes, we provide a computation of the throughputs for optimal routing policies under several information structures:

- Pre-allocation: the routing of a token is decided immediately upon entering the routing place, knowing the global marking.
- Allocation: the routing of a token can be decided at any instant, and knowing the global marking.
- Non-anticipative policy: the routing can be decided at any instant, knowing the global marking and the next transition available.

We compare the throughput that one can achieve using these different information structures, showing that the last one provides a better throughput than the second, which is also better than the first one.

We exhibit a free-choice net with one conflict place for which the optimal non-anticipative policy is to perform a race in some marking and a constant allocation in some other marking.

2 Stochastic free-choice nets

In this section, we recall the basic definitions of stochastic free-choice nets. We set $\mathbb{N}^* = \mathbb{N} \setminus \{0\}$ and $\mathbf{1}_X$ is the characteristic function of the set X .

A *Petri net* is a 4-tuple $\mathcal{N} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$ where $(\mathcal{P}, \mathcal{T}, \mathcal{F})$ is a directed bipartite graph with nodes $\mathcal{P} \cup \mathcal{T}$, $\mathcal{P} \cap \mathcal{T} = \emptyset$, and arcs $\mathcal{F} \subset (\mathcal{P} \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P})$ and where $M_0 \in \mathbb{N}^{\mathcal{P}}$. The elements of \mathcal{P} are called *places* and those of \mathcal{T} , *transitions*, and M_0 is called the *initial marking* of \mathcal{N} . For a node $x \in \mathcal{P} \cup \mathcal{T}$, we denote by $\bullet x$ the set of its predecessors and by x^\bullet the set of its successors. The marking evolves according to the *firing rule*: a transition a is *enabled* if: $\forall p \in \bullet a, M(p) \geq 1$. An

enabled transition can fire, and then the marking becomes M' with $M'(p) = M(p) - \mathbf{1}_{\bullet a}(p) + \mathbf{1}_{a\bullet}(p)$. This firing is denoted by $M \xrightarrow{a} M'$. A marking M' is *reachable from* M if there exists a sequence of transitions a_1, \dots, a_n such that M' is obtained from M by successively firing a_1, \dots, a_n . We write $M \xrightarrow{w} M'$ where $w = a_1 \cdots a_n$, and w is called an *admissible sequence*. We denote by $\mathcal{R}(M_0)$ the set of all the reachable markings (from M_0). For an admissible sequence $\sigma \in \mathcal{T}^*$, we denote by $\vec{\sigma}$ its *commutative image* (or Parikh vector), that is, the vector of $\mathbb{N}^{\mathcal{T}}$ that counts the number of occurrences of each transition in σ .

A *stochastic Petri net* is a Petri net where random timings have been added on the transitions. More precisely, a stochastic Petri net is a 5-tuple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0, \varphi)$, where $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M_0)$ is a Petri net, where $\varphi = (\varphi_a)_{a \in \mathcal{T}}$, and $\varphi_a = (\varphi_a(n))_{n \in \mathbb{N}^*}$ is a sequence of i.i.d. random variables with finite expectation ($E(\varphi_a(1)) < \infty$). Moreover, the sequences φ_a , $a \in \mathcal{T}$, are mutually independent. The firing rule is defined as follows: if the n -th firing of transition a starts at time t , then at time t , one token is removed from each input place of a , and at time $t + \varphi_a(n)$, one token is added in each output place of a .

A *free-choice (Petri) net* is a Petri net where: $\forall (p, a) \in \mathcal{P} \times \mathcal{T}$, $(p, a) \in \mathcal{F} \Rightarrow (p^\bullet = \{a\})$ or $(\bullet a = \{p\})$. That is, choices and synchronizations in the net are separated. A Petri net is *live* if for every reachable marking M' , and for every transition a , there exists a marking M'' reachable from M' such that a is enabled in M'' . A Petri net is *bounded* if there exists $m \in \mathbb{N}$ such that for every reachable marking $M \in \mathcal{R}(M_0)$, for every place $p \in \mathcal{P}$, $M(p) \leq m$. A connected live and bounded Petri net is strongly connected. In this article, we only consider strongly connected live and bounded free-choice nets.

The *cluster* $[x]$ of $x \in \mathcal{P} \cup \mathcal{T}$ is the smallest subset of $\mathcal{P} \cup \mathcal{T}$ such that: (i) $x \in [x]$; (ii) $p \in \mathcal{P}, p \in [x] \Rightarrow p^\bullet \in [x]$; (iii) $t \in \mathcal{T}, t \in [x] \Rightarrow \bullet t \in [x]$. The set of all the clusters of a Petri net defines a partition of the nodes. For free-choice Petri nets, each cluster contains only one place or only one transition.

Conflict resolution

In order to solve the conflicts in free-choice nets, at the places having several output transitions (conflict places), one needs to define a routing policy: when a token arrives in such a place, the policy defines which output transition will be fired with that token.

The *race policy* is defined as follows: when a token arrives in a conflict place p , every output transition of p begins its firing. The first transition that finishes to fire is effectively fired (it wins the race), and all the other output transitions of p abort their firing at that time. Therefore, the probability that a transition $a \in p^\bullet$ wins the race is $\mathbb{P}(\varphi_a(1) = \min_{a' \in p^\bullet} \{\varphi_{a'}(1)\})$, assuming that no ties are possible. (Otherwise a procedure to break ties needs to be specified.)

A *Bernoulli-routed* Petri net is a tuple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi, u)$ where $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi)$ is a stochastic Petri net and $u = (u_p)_{p \in \mathcal{P}}$ is the set of routing functions. For every place p , $u_p = (u_p(n))_{n \in \mathbb{N}^*}$ is a sequence of i.i.d. r.v.'s (hence the name Bernoulli routing), and those sequences are mutually independent and independent of the firing times. The r.v. $u_p(n)$ tells the transition that will be fired by the n -th token entering place p .

A Petri net with *periodic routing* is a tuple $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi, u)$ where $(\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi)$ is a stochastic Petri net and where $u = (u_p)_{p \in \mathcal{P}}$ with $u_p \in (p^\bullet)^{\mathbb{N}^*}$ being a deterministic periodic function. Again $u_p(n)$ tells the transition that will be fired by the n -th token entering place p .

A routing is *equitable* if for every conflict place p , each output transition is chosen with a strictly positive frequency. Under the race policy, the equitable condition becomes: for every place p , for every transition $a \in p^\bullet$,

$$\mathbb{P}(\varphi_a(1) = \min_{a' \in p^\bullet} \{\varphi_{a'}(1)\}) > 0. \quad (1)$$

In a general Petri net, where synchronizations and choices are not separated, the routing policy could lead to a deadlock (no transition can be fired) while the Petri net without routing is live. On the other hand, in the free-choice case, it is proved in [11] that every transition will fire infinitely often in a routed net if and only if the Petri net is live and the routings are equitable.

In the following, we always assume that equitability is satisfied.

3 Existence of the throughput

Theorem 1. *Let \mathcal{N}^k be a live and bounded stochastic free-choice net with a routing policy $k \in \{\text{race}, \text{Ber}, \text{per}\}$. For every transition b , there exists a constant $\lambda_b^k \in \mathbb{R}_+$ (throughput of transition b) such that a.s. and in L_1 ,*

$$\lim_{n \rightarrow \infty} \frac{n}{X_b^k(n)} = \lim_{t \rightarrow \infty} \frac{\mathcal{X}_b^k(t)}{t} = \lambda_b^k,$$

where $X_b^k(n)$ is the instant of completion of the n -th firing of transition b under policy k and $\mathcal{X}_b^k(t)$ is the number of firings completed at time t under policy k .

Proof. The result was proved in [11] for the Bernoulli routing.

-Race policy. The case of the race policy can be dealt with by showing that the behavior of the net under the race policy can be simulated by a suitable Bernoulli-routed net. Starting from $\mathcal{N}^{\text{race}} = (\mathcal{N}, \varphi)$, consider the Bernoulli-routed net $(\mathcal{N}, \varphi', u)$, where the distribution of transition a becomes

$$\mathbb{P}(\varphi'_a \leq t) = \mathbb{P}(\varphi_a \leq t | \forall b \in (\bullet a)^\bullet, \varphi_a \leq \varphi_b),$$

and where the routing function u is such that

$$\mathbb{P}(u_p(n) = a) = \mathbb{P}(\varphi_a \leq \varphi_b, \forall b \in (\bullet a)^\bullet).$$

(Here we assume for simplicity that $\mathbb{P}(\varphi_a = \varphi_b) = 0$ for all $a \neq b$.)

This Bernoulli routing is called the Bernoulli routing simulating the race policy. Indeed, in the race policy the output transition of a conflict place that is fired is that with the smallest firing time. The usual description is to begin the firing of every output transition and to fire the transition which is the first to finish its firing. Another equivalent way to proceed is to consider all the firing times of the output transitions before beginning to fire, compare them and choose the one with the smallest firing time. Under i.i.d. assumptions, the probability that transition a is fired is

$$\mathbb{P}(\varphi_a \leq \varphi_b, \forall b \in (\bullet a)^\bullet)$$

and the distribution of the firing time of a is

$$\mathbb{P}(\varphi'_a \leq t) = \mathbb{P}(\varphi_a \leq t | \forall b \in (\bullet a)^\bullet, \varphi_a \leq \varphi_b).$$

-Periodic routing. The case of periodic routings can be proved by adapting the proof of the Bernoulli case in [11].

Let \mathcal{E}_{per} be the state space of the net, that is

$$\mathcal{E}_{\text{per}} = \{(M, i_1, \dots, i_k) \mid M \in \mathcal{R}(M_0), i_j \in \{0, \dots, d_{p_j} - 1\}\}, \quad (2)$$

where p_1, \dots, p_k are the conflict places of the net. Clearly, this state space is finite.

Let block a transition b of the net. As the routing is equitable, there exists a unique reachable marking M_b , called the blocking marking of transition b , where transition b is the only enabled transition (see [11]). Let the state obtained be (M_b, r) . If b is fired and blocked again, the new state obtained is (M_b, r') . Moreover, r' does not depend on the order of the transition fires (because the routing is deterministic), so (M_b, r) is the only state that can be reached ([11], Lemma 4.4). Up to a commutation of transitions, the firing sequence is also unique.

If transition b is successively blocked and fired, according to the previous description, there exists a state (M_b, r) that is visited infinitely often. Furthermore, the firing sequence between two visits of that state is unique up to the commutation of transitions of the firing sequence. Let k_t be the number of occurrences of t in such a sequence.

We suppose now, without loss of generality, that the initial state of the net is (M_b, r) , and denote by $(\tilde{\mathcal{N}}, \tilde{M}_b, \tilde{\tau}, \tilde{u})$ the open expansion of $(\mathcal{N}, M_b, \tau, u)$ and b defined as follows.

Let $K = \max\{k : M_b \xrightarrow{b^k}\}$. The open expansion of $(\mathcal{N}, M_b, \tau, u)$ and b is:

- $\tilde{\mathcal{P}} = \mathcal{P} \cup \{p_b, p_I\}$;
- $\tilde{\mathcal{T}} = (\mathcal{T} - \{b\}) \cup \{I, b_i, b_o\}$;
- $\tilde{\mathcal{F}} = \begin{cases} (\mathcal{F} - \{(p, b) \in \mathcal{F}, (b, p) \in \mathcal{F}\}) \\ \cup \{(p, b_o) : (p, b) \in \mathcal{F}, (b, p) \notin \mathcal{F}\} \\ \cup \{(b_i, p) : (b, p) \in \mathcal{F}, (p, b) \notin \mathcal{F}\} \\ \cup \{(b_i, p), (p, b_i) : (b, p) \in \mathcal{F}, (p, b) \in \mathcal{F}\} \\ \cup \{(I, p_I), (p_I, b_i), (b_o, p_b), (p_b, b_i)\} \end{cases}$;
- $\tilde{M}_p = \begin{cases} M_p & \text{if } p \in \mathcal{P} - \bullet b \\ M_p - K + K\chi_{b\bullet} & \text{if } p \in \bullet b \\ K & \text{if } p = p_b \\ 0 & \text{if } p = p_I \end{cases}$;
- $\tilde{\tau}_a(n) = \begin{cases} \tau_a(n) & \text{if } a \in \mathcal{T} - \{b\} \\ \tau_b(n) & \text{if } a = b_i \\ 0 & \text{if } a = b_o \end{cases}$,
- $\tilde{u}_p(n) = u_p(n)$.




Figure 1: Open expansion.

Figure 1 illustrates this construction. The Petri net obtained is neither live nor bounded. In order to be enable some transitions, one has to define the firing times of transition I . When transition I is saturated (transition I is fired an infinitely number of time at instant 0), \mathcal{N} and $\tilde{\mathcal{N}}$ have the same behavior. We now prove the existence of the throughput of transition b in $\tilde{\mathcal{N}}$. Let define ζ_1 as the random vector $(\tau_1(1), \dots, \tau_1(k_1), \dots, \tau_{|\mathcal{T}|}(1), \dots, \tau_{|\mathcal{T}|}(k_{|\mathcal{T}|}))$ and more generally,

$$\zeta_i(\tau_1(((i-1)k_1 + 1), \dots, \tau_1(ik_1), \dots, \tau_{|\mathcal{T}|}(((i-1)k_{|\mathcal{T}|} + 1), \dots, \tau_{|\mathcal{T}|}(ik_{|\mathcal{T}|})).$$

The vector ζ_i is the sequence of the firing times of the transitions between the i -th and the $i+1$ -th visit in state (M_b, r) if the firing dates of I are spaced enough so that transition b is fired exactly k_b times. Because of the independence hypothesis, the sequence (ζ_i) is i.i.d. The existence of the throughput is now a consequence of the monotone-separable framework defined in [4] that can be applied here. \square

Remark The transformation of a Petri net with the race policy into the Bernoulli routing simulating the race policy modifies the distributions if the firing times. This leads to more complex distributions. In particular, a net with exponentially distributed firing times for the race policy is not transformed into a Bernoulli routed net with exponential firing times.

3.1 Ratio between the throughputs of the transitions

Although it seems impossible to compute the throughput of the transitions when the firings have general distributions, it is rather easy to compute the ratio between the throughputs of two different transitions for all three routing policies.

Define the *routing matrix* $R^k = (R_{ij}^k)_{i,j \in \mathcal{T}}$ as:

$$R_{i,j}^k = \frac{1}{|\bullet j|} \sum_{p \in \mathcal{P}: i \rightarrow p \rightarrow j} F^k(p, j),$$

where $F^k(p, j)$ is the frequency of routing to transition j from place p under the routing k . In particular, $F^{Ber}(p, j) = \mathbb{P}(u_p(1) = j)$, $F^{race}(p, j) = \mathbb{P}(\varphi_j \leq \varphi_a, \forall a \in p^\bullet)$, and $F^{per}(p, j)$ is the proportion of tokens routed to j over one period of the routing, that is, $|r_p|_j/d_p$, where d_p is the period on the routing in place p and $|r_p|_j$ the number of times transition j is chosen during a period.

From the equitable assumption, in all three cases, the matrix R^k is irreducible, its spectral radius is 1, and it admits a unique eigenvector $x^k = (x_a^k)_{a \in \mathcal{T}}$, $x_a^k \in \mathbb{R}_+ \setminus \{0\}$, $\sum_a x_a^k = 1$, such that $x^k R^k = x^k$.

Theorem 2. *The model is the same as in Theorem 1. For all routing policy k belonging to $\{race, Ber, per\}$, there exists a constant $c^k \in \mathbb{R}_+ \cup \{\infty\}$ such that for all transition a , $\lambda_a^k = c^k x_a^k$.*

The proof is an adaptation of the proof of [11, Prop. 5.1] in which Bernoulli-routed nets are considered.

Proof. As the result has already been proved for Bernoulli routed nets, the result holds for the race policy, using the transformation of a net with the race policy into a Bernoulli-routed net. We now show the theorem for periodic routings.

We first suppose that λ_a is finite for every transition $a \in \mathcal{T}$. Let a be a transition and $\mathcal{X}_a(t)$ be the number of complete fires at time t . Let $\mathcal{Y}_{pa}(t)$, $a \in \mathcal{T}$, $p \in \bullet a$, be the number of tokens routed in place p to transition a at time t . Then,

$$\mathcal{X}_a(t) \leq \mathcal{Y}_{pa}(t) \leq \mathcal{X}_a(t) + \bar{M}, \quad (3)$$

where \bar{M} is the bound of \mathcal{N} . As the equation

$$\mathcal{Y}_{pa}(t) = \sum_{i=1}^{K(t)} \mathbf{1}_{\{u_p(i)=a\}}, K(t) = M_0(p) + \sum_{b \in \bullet p} \mathcal{X}_b(t) \quad (4)$$

also holds, equation (3) gives

$$\lambda_a = \lim_t \frac{\mathcal{X}_a(t)}{t} = \lim_t \frac{\mathcal{Y}_{pa}(t)}{t}, \quad (5)$$

and equation (4), using the periodicity of u_p , gives

$$\lim_t \frac{\mathcal{Y}_{pa}(t)}{t} = \lim_t \frac{\sum_{i=1}^{K(t)} \mathbf{1}_{\{u_p(i)=a\}}}{K(t)} \cdot \frac{K(t)}{t} = \frac{|r_p|_a}{d_p} \cdot \frac{M_0(p) + \sum_{b \in \bullet p} \mathcal{X}_b(t)}{t}. \quad (6)$$

As the throughput exists, we get

$$\lambda_a = \frac{|r_p|_a}{d_p} \sum_{b \in \bullet p} \lambda_b. \quad (7)$$

Moreover, this equation is valid for every place $p \in \bullet a$. So

$$\lambda_a = \frac{1}{|\bullet a|} \sum_{p \in \bullet a} \frac{|r_p|_a}{d_p} \sum_{b \in \bullet p} \lambda_b = \sum_{b \in \bullet p} \lambda_b R_{ba} \quad (8)$$

and $\lambda = \lambda R$.

By construction, R is non-negative. As the routing is equitable, we have $\forall (p, j) \in (\mathcal{P} \times \mathcal{T}) \cap \mathcal{F}$, $|r_p|_a/d_p > 0$, so for every $i, j \in \mathcal{T}$ such that there exists $p \in \mathcal{P}$ such that $i \rightarrow p \rightarrow j$, $R_{ij} > 0$. Moreover, \mathcal{N} is strongly connected, so R is irreducible. Using the Perron-Frobenius theorem, we deduce that λ is unique (up to a multiplicative factor) and positive. The spectral radius of R is then 1.

If there exists a transition a such that $\lambda_a = \infty$, then $\lambda = (\infty, \dots, \infty)$ because R is irreducible. Moreover, R depends only on the routing frequencies of the tokens and not on the timings. So, if every timing is replaced by an exponentially distributed r.a. of parameter 1, and the routing is not modified, then, the timing of each fire is almost surely positive and $\lambda \in \mathbb{R}_+$. Applying the first part of the proof, the spectral radius of R is 1. \square

Observe that R depends only on the routing frequencies of tokens, and not on the timings of the transitions. Therefore, the ratios are the same for the three policies provided that $F^k(p, j)$ are equal for all three policies.

3.2 Comparison between the policies

As mentioned before, the ratio between the throughputs of the transitions are the same in all three cases. Therefore, to compare the throughputs, one just has to compare the total throughputs $c^k = \sum_{a \in \mathcal{T}} \lambda_a^k$, $k \in \{\text{race}, \text{Ber}, \text{per}\}$.

Proposition 1. *Let $\mathcal{N}^{\text{race}} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi)$ be a live and bounded stochastic free-choice net with the race policy and let $\mathcal{N}^{\text{Ber}} = (\mathcal{P}, \mathcal{T}, \mathcal{F}, M, \varphi, u)$ be the Bernoulli-routed net with the same firing times and routing frequencies as $\mathcal{N}^{\text{race}}$. Then, $c^{\text{race}} \geq c^{\text{Ber}}$.*

Let first prove a preliminary lemma.

Lemma 1. *Let p be a conflict place with output transitions $p^\bullet = \{a_1, \dots, a_k\}$. For $t \in \mathbb{R}_+$, let $P_B(t) = \mathbb{P}(\varphi_{a_1} < t)$ be the probability that the firing time of a_1 in \mathcal{N}^{Ber} is less than t and $P_r(t) = \mathbb{P}(\varphi_{a_1} < t | \varphi_{a_1} < \varphi_{a_i}, i \in \{2, \dots, k\})$ the probability that the firing time of a_1 in $\mathcal{N}^{\text{race}}$ is less than t . Then, $\forall t \in \mathbb{R}_+$, $P_B(t) \leq P_r(t)$.*

Proof. Let $P_1(t) = \mathbb{P}(\varphi_{a_1} < t)$ and $P_2(t) = \mathbb{P}(\min(\varphi_{a_2}, \dots, \varphi_{a_k}) < t)$. We have $P_B(t) = \int_0^t dP_1(u)$ and

$$\begin{aligned} P_r(t) &= \frac{\mathbb{P}(\varphi_{a_1} < t \wedge \varphi_{a_1} < \varphi_{a_i}, i \in \{2, \dots, k\})}{\mathbb{P}(\varphi_{a_1} < \varphi_{a_i}, i \in \{2, \dots, k\})} \\ &= P_r(t) = \frac{\int_0^\infty \int_0^{t \wedge u} dP_1(v) dP_2(u)}{\int_0^\infty \int_0^u dP_1(v) dP_2(u)}. \end{aligned}$$

But is

$$\begin{aligned} P_B(t) \cdot \int_0^\infty \int_0^u dP_1(v) dP_2(u) &= \int_0^\infty \int_0^u dP_1(v) dP_2(u) \times \int_0^t dP_1(v) \\ &= \int_0^\infty \left[\int_0^u dP_1(v) \right] \left[\int_0^t dP_1(v) \right] dP_2(u) \\ &\leq \int_0^\infty \int_0^{t \wedge u} dP_1(v) dP_2(u) \\ &= P_r(t) \cdot \int_0^\infty \int_0^u dP_1(v) dP_2(u), \end{aligned}$$

because

$$\int_0^{t \wedge u} dP_1(v) = \min\left(\int_0^t dP_1(v), \int_0^u dP_1(v)\right)$$

and

$$\int_0^x dP_1(v) \leq 1 \quad \forall x.$$

□

That last lemma can be interpreted in terms of *stochastic comparison*, notion that has been developed in [13]: the random variable $\varphi_{a,r}$ is stochastically less than $\varphi_{a,B}$ ($\varphi_{a,r} \leq_{st} \varphi_{a,B}$). Now, we can prove Proposition 1.

Proof of Proposition 1. For every transition $a \in \mathcal{T}$, $X_a(n)$ is the completion time of the n -th fire of transition a . Let $\nu_{ap} = \min\{k : \sum_{i=1}^k \mathbf{1}_{\{u_p(i)=a\}} = n\}$ be the minimal number of tokens arrived in p such that n tokens have been routed to a . Then $X_a(n)$ satisfies the equation ([3])

$$X_a(n) = \left\{ \max_{p \in \bullet a} \left[\min_{\substack{(n_i, i \in \bullet p): M_p + \\ \sum_{i \in \bullet p} n_i = \nu_{pa}(n)}} \max_{i \in \bullet p} X_i(n_i) \right] \right\} + \varphi_a(n),$$

where $\varphi_a(n)$ is the firing time of the n -th fire of transition a (if \mathcal{N}^{race} is considered, we will write $\varphi_{a,r}$ and if \mathcal{N}^{Ber} is considered, we write $\varphi_{a,B}$). That equation is satisfied by both \mathcal{N}^{race} and \mathcal{N}^{Ber} . When those equation are unfolded, they can be expressed in function of $\varphi_b(i)$, $b \in \mathcal{T}$ only.

As

$$\Phi : ((x_{i,n_i})_i, y) \rightarrow \left\{ \max_{p \in \bullet a} \left[\min_{\substack{(n_i, i \in \bullet p): M_p + \\ \sum_{i \in \bullet p} n_i = \nu_{pa}(n)}} \max_{i \in \bullet p} x_{i,n_i} \right] \right\} + y$$

is non-decreasing, X_a is non-decreasing. Moreover, to compute $X_a(n)$, only a finite number of $\varphi_b(i)$ is necessary, which depends only on the routing. Indeed, let M_a be the marking where only a a can fire (that marking exists and is unique in routed free-choice nets, see [11, Lemma 4.4]), let $\vec{\sigma}$ be the commutative image (unique for a given routing) of the fired transitions from M to M_a after $n-1$ firings of a . Then $X_a(n)$ depends only on $\varphi_b(1), \dots, \varphi_b(|\vec{\sigma}|_b)$, $b \in \mathcal{T}$, as the other firing can only happen after the marking of transition a .

Theorem 3.3.11 in [13] states that the stochastic order is stable for non-decreasing functions. From Lemma 1, for every transition b , for every $n \in \mathbb{N}^*$, $\varphi_{r,b}(n) \leq_{st} \varphi_{B,b}(n)$ and as Φ is non-decreasing, we get $\lambda_a^{race} = \lim_{n \rightarrow \infty} \frac{n}{X_{r,a}(n)} \geq \lim_{n \rightarrow \infty} \frac{n}{X_{B,a}(n)} = \lambda_a^{Ber}$ and $c^{race} \geq c^{Ber}$. □

The comparison with periodic routings is more difficult. This will be illustrated in the next section which focuses on computational issues.

4 Computing Throughputs

This section is devoted to the computation of the throughput in live and bounded free-choice Petri nets. We now consider that every transition a has a firing time exponentially distributed with parameter $\mu_a \in (0, \infty)$.

4.1 Race policy

The race policy case is well-known. The marking evolves as a *continuous-time jump Markov process*. Let M be a reachable marking and T_M be the set of the transitions enabled at M . The first transition fired is $a \in T_M$ with probability $\mu_a / (\sum_{a' \in T_M} \mu_{a'})$. The firing time is exponentially distributed with parameter $\sum_{a' \in T_M} \mu_{a'}$. The stationary distribution π_r of this process is characterized by the equation $\pi_r Q = 0$, where Q is the infinitesimal generator defined as follows. Denote




Figure 2: Example of a live and bounded free-choice Petri net.

by $M \cdot a$ the marking such that $M \xrightarrow{a} M \cdot a$. We have

$$\begin{aligned} \forall M_1 \in \mathcal{R}(M_0), \quad Q_{M_1, M_1 \cdot a} &= \mu_a \text{ if } a \in T_{M_1} \\ Q_{M_1, M_1} &= - \sum_{a \in T_{M_1}} Q_{M_1, M_1 \cdot a} \end{aligned}$$

The total throughput is then given by the formula:

$$c^{race} = \sum_{M_1 \in \mathcal{R}(M_0)} -(\pi_r)_{M_1} \cdot Q_{M_1, M_1}.$$

To illustrate the computation of the throughput, we study an example that will be used throughout the paper.

Example 1. Consider the Petri net in Figure 2. The places are named by letters (a to e) and the transitions by numbers (1 to 5). The parameters of the exponentially distributed timings of the transitions are respectively $\mu_1 = 2$, $\mu_2 = 2$, $\mu_3 = 3$, $\mu_4 = 5$, and $\mu_5 = 1$. The state space is the set of reachable markings: $\mathcal{R} = \{\{a, d\}, \{a, e\}, \{b, d\}, \{b, e\}, \{c, d\}, \{c, e\}\}$. The marking is a continuous-time Markov process with infinitesimal generator:

$$Q = \begin{pmatrix} -2 & 0 & 2 & 0 & 0 & 0 \\ 1 & -3 & 0 & 2 & 0 & 0 \\ 2 & 0 & -5 & 0 & 3 & 0 \\ 0 & 2 & 1 & -6 & 0 & 3 \\ 0 & 5 & 0 & 0 & -5 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{pmatrix}.$$

The stationary distribution π^{race} is obtained by solving the equation $\pi^{race} Q = 0$. We obtain $\pi^{race} = (85, 90, 40, 30, 42, 90)/337$.

The throughput is $c^{race} = \sum_i -\pi_i \cdot Q_{ii} = 1120/377 \approx 2.97$.

4.2 Bernoulli routings

For a free-choice Petri net with Bernoulli routings, the marking is not a Markov process anymore. One possibility is to add immediate firing transitions to model the routing which would yield a semi-Markov process for the marking (see [1]). Another possibility, used here, is to model the evolution of the net by a Markov process over an extended state space. This approach has the advantage that computations can be carried out symbolically which is very helpful for optimization purposes.

The main trick in the construction is the choice of the state space. When a token enters a conflict place, the timing of the transition to be fired, depends on the routing. If we took $\mathcal{R}(M_0)$ for the state space, then this would lead to difficulties due to that dependence. In order to separate the timings from the routing, we consider a new state space: when a token enters a choice place, the transition it can fire is already defined. For $M \in \mathcal{R}(M_0)$, let $\mathcal{T}(M)$ be the set of all the maximal sets of transitions that can be fired simultaneously (in the non-timed Petri net) under M . Then, the extended state space is

$$\mathcal{E} = \{(M, T) \mid M \in \mathcal{R}(M_0), T \in \mathcal{T}(M)\}, \quad (9)$$

i.e. every state corresponds to a pair formed by a marking and a set of enabled transitions.

The infinitesimal generator Q of the chain is defined as follows. Let $(M_1, T_1) \in \mathcal{E}$, and $a \in T_1$. Transition a is fired with rate μ_a and the new set of enabled transitions is $T_2 = (T_1 \setminus \{a\}) \cup T'$ where T' is a maximal set of newly enabled transitions, chosen randomly according to the Bernoulli routings.

Example 2. Consider again the Petri net of Figure 2. A token arriving in place b fires transition 2 with probability p and transition 3 with probability $q = 1 - p$. The state space is $\{(ad, 1), (ae, 1, 5), (bd, 2), (bd, 3), (be, 2, 5), (be, 3, 5), (cd, 4), (ce, 5)\}$, and the infinitesimal generator is :

$$Q = \begin{pmatrix} -2 & 0 & 2p & 2q & 0 & 0 & 0 & 0 \\ 1 & -3 & 0 & 0 & 2p & 2q & 0 & 0 \\ 2 & 0 & -2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -3 & 0 & 0 & 3 & 0 \\ 0 & 2 & 1 & 0 & -3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & -4 & 0 & 3 \\ 0 & 5 & 0 & 0 & 0 & 0 & -5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{pmatrix}.$$

By solving $\pi^{Ber} Q = 0$, we get π^{Ber} formally, each coordinate being a rational fraction of p . The total throughput is :

$$c^{Ber} = \frac{60(4p^2 - 17p + 18)}{138p^2 - 403p + 414}.$$

The maximum of c^{Ber} is reached for $p = (846 - 30\sqrt{615})/751 \approx 0.14$. The corresponding value of the throughput is approximatively 2.61.

To have the same routing probabilities as in the race policy case, one must take $p = 2/5$. The stationary probability is then $\pi^{Ber} = (190, 270, 100, 141, 72, 108, 222, 81)/1184$, and the throughput is $c^{Ber} = 1480/573 \approx 2.58$.

In both cases, we computed the total throughputs c^{race} and c^{Ber} . To get back to the throughput of one transition, we also need to compute the left-eigenvector associated to the eigenvalue 1 of the matrix R . We have

$$R = \begin{pmatrix} 0 & p & 1-p & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 2/3 & 0 & 0 & 0 & 1/3 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix},$$

and then, with $p = 2/5$, $\lambda^k = \frac{c^k}{20}(10, 4, 2, 3, 1)$, where $k \in \{Ber, race\}$.

In the above example, the maximum of c^{Ber} is strictly less than c^{race} . This is not always the case and it is easy to build models in which a far better throughput can be reached with Bernoulli routings than with the race policy.

4.3 Periodic routings

We assume that place p_i has a periodic routing policy with period d_i . The behavior of the net can be modelled by a continuous-time Markov process with a state space $E = \mathcal{R}(M_0) \times \{0, \dots, d_1 - 1\} \times \dots \times \{0, \dots, d_s - 1\}$ where $s = |\mathcal{P}|$.

Being in a state (M, r_1, \dots, r_s) means that the current marking is M and that the next transition to be chosen by a token in place p_i is given by the r_i -th element in the periodic sequence attached to i . The infinitesimal generator is defined accordingly.

The number of states becomes rapidly large when the periods of the routing functions increase. Some numerical computations have been carried out using Maple for the foregoing example of Figure 2. For every periodic routing of the conflict place b that has period at most 10, the transition matrix has been automatically generated and then the throughput computed. The results are displayed in Figure 3. For a given ration, different values are computed and correspond to the different throughputs obtained with different routings with the same proportion of routings.




Figure 3: Throughput of several periodic routing functions when the proportion of tokens sent to transition 2 varies.

Numerical evidence suggests that the best periodic routing is given by *balanced routing functions*. This fact has been proven for many open systems (see [2]). In the case of closed systems (as for the free-choice net used here), this is in general unproved, with a few exceptions, see [8].

Numerical evidence also suggests that the best periodic routing is better than the best Bernoulli routing, see Figure 4. But, the throughput obtained for an arbitrary periodic routing can be less than the throughput obtained with Bernoulli routing with the same proportions.

The maximal throughput for the periodic routing is reached when the proportion of tokens sent to transition 2 is 0.5. This is in contrast with the situation of Bernoulli routings, where we recall that the maximum was attained for a proportion approximately equal to 0.14. This means that one can't find the best periodic routing by first computing the best Bernoulli routing and then taking the balanced periodic routing corresponding to the same routing proportions. However, in the next section we show how to find the best the Bernoulli routing.

The shape of the curve for the best periodic routing in Figure 4 is characteristic. It seems to be piecewise-affine with singularities at rational points with small denominators. This is reminiscent of the numerical data obtained in [8] for a closed free-choice Petri net with deterministic timings.




Figure 4: Comparison of the Bernoulli routing with the best periodic routing when the proportion of tokens sent to transition 2 varies.

5 Optimal routing policies

5.1 Optimal Bernoulli routing

In the previous sections, we showed the existence of the throughput in free-choice nets with the race policy, and proved that the throughput obtained with that policy is greater than with the routed net with the same routing probabilities. In this section, we focus on the routed nets and study the dependence of the throughput with the routing parameters. In particular we provide an efficient method to optimize the Bernoulli routing parameters in one conflict place in order to maximize the total throughput of the net.

5.1.1 Optimization of one parameter

We first suppose that every routing parameter is fixed except one, say p . The average time spent in a state does not depend on the Bernoulli parameters, so only the stationary distribution of the embedded Markov chain of the Markov process described above depends on p .

Let $P(p)$ be the transition matrix of the embedded Markov chain. The coefficients of $P(p)$ are polynomial of degree at most 1. For every $p \in]0, 1[$, it is a stochastic and irreducible matrix, so, by the Perron-Frobenius theorem, $P(p) - I$ has rank $r - 1$, where r is the dimension of the matrix $P(p)$. The stationary probability $\pi(p)$ is a solution of the equation

$$x.(P(p) - I) = 0, \quad (10)$$

and every solution of that equation is collinear to $\pi(p)$.

The following theorem ([12], Theorem 6.1) gives the complexity of a recursive algorithm that computes a polynomial solution of $x.(P(p) - I) = 0$ in cubic time.

Theorem 3. *Let $M \in F[x]^{n \times m}$ be of rank r and degree bounded by d . There exists an algorithm solving the system $x.M = 0$ in $O(nmrd^2)$ elementary operations.*

Applied to our case, if n is the size of the state space defined in (9), the cost for solving $x(P(p) - I) = 0$ is $O(n^2(n - 1)) = O(n^3)$. Note that the state space can be exponentially large




Figure 5: Example of a Petri net where the maximum throughput with Bernoulli routing can be higher than with the race policy.

compared to the size of the net: it is larger than the reachable marking set. The solution given by the algorithm is a polynomial vector, whose degree is bounded by n . In fact, it can be shown rather easily that the maximal degree of the coefficients of w is the number of lines where the parameter p appears in $P(p)$.

Notice that the time complexity to compute $\pi(p)$ is the same as the time complexity to compute the solution of a linear system with real coefficients. A more naive method would have led to a time complexity in $O(n^6)$: as it is possible to bound the degree of a polynomial solution x of Equation 10, we can write $x(p)$ as $x_0 + x_1.p + \dots + x_d.p^d$, where d is the maximal degree of $x(p)$ and x_i are real vectors. Now, $P(p)$ can be written as $P_0 + P_1.p$, where P_0 and P_1 are real matrices. By identifying the coefficients, the system can now be solved as a linear system with $d.n$ equations, which leads to a time complexity in $O((n.d)^3)$.

The stationary probability is the normalized vector of that solution: it is a rational fraction, with degree at most n for the numerator and denominator, and the throughput of the net is a linear function of $\pi(p)$. Then, maximizing P consists in studying a rational fraction between 0 and 1, and finding its maximum.

Example 3. *Let us consider the net displayed in Figure 2. The total throughput as well as the stationary measure π of the Markov chain can be computed as a rational function of p . The maximum degree of the numerator of π is at most the number of lines in which the parameter appears. We find:*

$$\pi^{Ber}(p) = \frac{1}{4p^2 - 17p + 18} \left(\frac{2p + 3}{2}, \frac{9(p - 1)}{2}, \frac{5p}{2}, \frac{9 - p^2 - 8p}{4}, 3p(1 - p), 3(1 - p)^2, \frac{4p^2 - 13p + 9}{2}, \frac{9(p^2 - 2p + 1)}{4} \right),$$

which gives the throughput: $c^{Ber}(p) = \frac{60(4p^2 - 17p + 18)}{139p^2 - 403p + 414}$. Then, we get that the maximum is reached for $p = \frac{846 - 30\sqrt{615}}{751} \approx 0.14$ and the maximum throughput is approximately 2.61.

Note that in this example, the total throughput achieved with all Bernoulli routings (*i.e.* with all possible routing parameters) is always less than the throughput obtained with the race policy. This is not always the case. The Petri net represented in Figure 5 (with exponential parameters of the timings given in bold in the figure) can achieve a far better throughput with Bernoulli routings than with the race policy. With the race policy, $c^{race} \approx 2.07$. If the Bernoulli routing parameters in places c and e , have same proportions as for the race policy and if p is the probability of routings tokens in place a towards transition 1, the throughput is $c^{ber}(p) = \frac{2300}{1400 - 1089p}$. The maximum is reached when p is close to 1 and the total throughput is $c^{Ber}(1) \approx 7.40$.

5.1.2 General case

When there are more than one parameter, there is no known efficient algorithm to compute the throughput. But the naive method with several parameters is still valid: there is a solution of (10) such that the degree of each parameter is less than n (because the maximal degree of the coefficients of $P(p)$ is at most 1). By identifying the coefficients of same degree in the equation, it can be solved as a linear system of $n(n+1)^m$ equations, where m is the number of parameters. Then the complexity is $O(n^{3m+3})$, which could be too large for practical computations as soon as the state space increases.

5.2 Optimal allocations

We now consider larger classes of routings. In this section, we consider the *optimal* routing policies in a free-choice net with respect to the throughput. Again, firing times will be exponential, so that an MDP (Markov Decision Process) approach will be possible. The optimal policy under several information structures, are derived and the corresponding throughputs are computed and compared. This provides an enlightening illustration of the value of information.

Now, consider a stochastic live and bounded free-choice net with exponential firing times, the firing rate of transition a being $\mu_a \in (0, \infty)$. It is convenient to view the model as follows. The instants of *potential firing* or *availability* of a given transition a are given by an exogenous Poisson process of rate μ_a . At a potential firing instant, a firing may occur if and only if the transition is enabled. In this last case, firing or not the transition is the role of a *decision maker*. The goal of the decision maker is to maximize the total throughput. Variations of the model are obtained depending on the moment when the decisions need to be taken and the quantity of information available.

To use a MDP approach, it is convenient to work in a discrete-time setting. To that purpose, the process obtained by superposition of the Poisson processes of rate μ_a is replaced by a sequence of i.i.d. r.v.'s valued in \mathcal{T} (the set of transitions) and of distribution $(\mu_a/\Lambda)_{a \in \mathcal{T}}$, where $\Lambda = \sum_a \mu_a$. Now, time is slotted, and at each time slot, precisely one transition has the potential to fire. The firing will occur, at this same time slot, if: (i) the transition is enabled, (ii) the decision maker agrees. The above is a simple instance of the standard “uniformization” trick.

The immediate reward at each slot is 1 if a transition is fired and is 0 otherwise. Maximizing the throughput is now equivalent to maximizing the infinite horizon average reward. Therefore, it is possible to model the maximizing problem using a MDP. The maximal throughput and the optimal policy will be given by the Bellman equation associated with the MDP (see for example [14]). In particular, the maximal throughput of the Petri net is the average reward per unit of time of the MDP multiplied by the uniformization constant Λ . Here, the state space of the MDP is always finite. So the Bellman equation can be explicitly solved using policy iteration.

5.2.1 Optimal token pre-allocation

Assume that a token enters a conflict place at time slot n . The decision maker has to choose, immediately, one of the output transitions. The information available is the marking of the Petri net at time slot n . The token will eventually fire the chosen transition at the first slot after n when it becomes available. In particular, when the decision is taken, it is not known which one of the output transitions will be available first. We call this a *pre-allocation* policy.

Bernoulli routings and periodic routings are special cases of pre-allocations where the knowledge of the global marking is not used.

The state space of the MDP is formed by the set of all pairs formed by a reachable marking (M) and an allocation of tokens in conflict places (ρ), which must be memorized as the allocation in a conflict place does not change until the token is consumed.

Let $M \cdot t$ be the marking obtained from M by the firing of transition t . For $t \in \rho$, let $\text{rout}(M, \rho, M \cdot t)$ be the set of possible allocations for the new tokens appearing in conflict places when the firing $M \xrightarrow{t} M \cdot t$ is performed. Let $\text{rout}(M, \rho) = \prod_{t \in \rho} \text{rout}(M, \rho, M \cdot t)$ be the set of

all possible future decisions in state (M, ρ) . For $r \in \text{rout}(M, \rho)$ and $t \in \rho$, the firing of t will transform the state (M, ρ) into the state $(M \cdot t, \rho \cdot (t, r))$ where $\rho \cdot (t, r) = \rho \setminus \{t\} \cup (t^{\bullet\bullet} \cap r(t))$. In other words, the allocation remains the same except for the newly enabled places. As far as conflict places are concerned, one and only one output transition of each newly marked place is chosen.

For a routing decision $r \in \text{rout}(M, \rho)$, the expected gain at step k is

$$J^k(M, \rho) = \sum_{t \in \rho} \left[\frac{\mu_t}{\Lambda} (J^{k+1}(M \cdot t, \rho \cdot (t, r)) + 1) + \frac{\Lambda - \sum_{t \in \rho} \mu_t}{\Lambda} J^{k+1}(M, \rho) \right].$$

Then the optimal policy is to maximize the expectation:

$$J^k(M, \rho) = \max_{r \in \text{rout}(M, \rho)} \sum_{t \in \rho} \left[\frac{\mu_t}{\Lambda} (J^{k+1}(M \cdot t, \rho \cdot (t, r)) + 1) + \frac{\Lambda - \sum_{t \in \rho} \mu_t}{\Lambda} J^{k+1}(M, \rho) \right].$$

This Markovian Decision process is finite and irreducible, therefore, using Proposition 2.1, Chap 4. Vol. II in [5], there exists a unique reward vector, $J = (J(M, \rho))_{(M, \rho)}$ and a unique optimal average reward, g verifying the Bellman equation corresponding to that process:

$$J(M, \rho) + g = \max_{r \in \text{rout}(M, \rho)} \left(\sum_{t \in \rho} \frac{\mu_t}{\Lambda} (J(M \cdot t, \rho \cdot (t, r)) + 1) + \frac{\Lambda - \sum_{t \in \rho} \mu_t}{\Lambda} J(M, \rho) \right).$$

Example 4. Consider the example of Figure 2. Solving the Bellman equation, we get the following optimal pre-allocation. When a token enters place b , allocate it to:

- Transition 2 if the current marking is $\{b, e\}$;
- Transition 3 if the current marking is $\{b, d\}$.

The corresponding throughput is approximately 2.8.

In order to solve the Bellman equation, we used dynamic programming techniques and the algorithm of values iteration.

5.2.2 Optimal token allocation

Each token in a conflict place is allocated to an output transition, but this allocation can be modified by the decision maker at the beginning of each time slot: (i) knowing the current marking, (ii) but not knowing the transition which is about to become available. We call this a (*token*) allocation policy. Pre-allocation is of course a special case of allocation policy.

The state space of the MDP is simply the set of reachable markings. Let $\mathcal{T}(M)$ be the set of all maximal sets of transitions that can fire in marking M . Each element of $\mathcal{T}(M)$ contains exactly one transition in each marked cluster. The new Bellman equation is:

$$J(M) + g = \max_{r \in \mathcal{T}(M)} \left(\sum_{t \in r} \frac{\mu_t}{\Lambda} (J(M \cdot t) + 1) + \frac{\Lambda - \sum_{t \in r} \mu_t}{\Lambda} J(M) \right).$$

Example 5. Consider the model of Figure 2. We get the following optimal allocation. At the beginning of a time slot, if there is a token in place b , allocate it to:

- Transition 2 if the current marking is $\{b, e\}$;
- Transition 3 if the current marking is $\{b, d\}$.

This optimal policy is different from the one in Example 4. Assume that the marking is $\{b, e\}$. Then the token in b is allocated to transition 2, but if transition 5 fires first, then the token in b gets re-allocated to transition 3. The corresponding throughput is approximately 3.05.

5.2.3 Optimal non-anticipative policy

At each time slot, if the available transition is enabled, the decision maker decides either to fire or not to fire the transition. The available information is the current marking. This can also be viewed as a model where the decision maker may reallocate the token at the beginning of each time slot knowing: (i) the current marking, (ii) the transition which is about to become available. This is the maximal amount of information available without look-ahead. So we call this a *non-anticipative routing* policy.

Clearly token allocation policies form a subset of non-anticipative policies. But the race policy can also be emulated by a non-anticipative policy. It is also possible to have a coexistence of allocations and races, see Example 6.

The state space is still the reachability graph. But the set of possible decisions in marking M is now $\mathcal{E}(M)$, where $\mathcal{E}(M)$ is the set of all subsets of \mathcal{T} containing *at least* one transition in each marked cluster. Observe that the set $\mathcal{E}(M)$ is larger than the set $\mathcal{T}(M)$ of possible decisions for token allocation policies. The new Bellman equation is

$$J(M) + g = \max_{r \in \mathcal{E}(M)} \left(\sum_{t \in r} \frac{\mu_t}{\Lambda} (J(M \cdot t) + 1) + \frac{\Lambda - \sum_{t \in r} \mu_t}{\Lambda} J(M) \right).$$

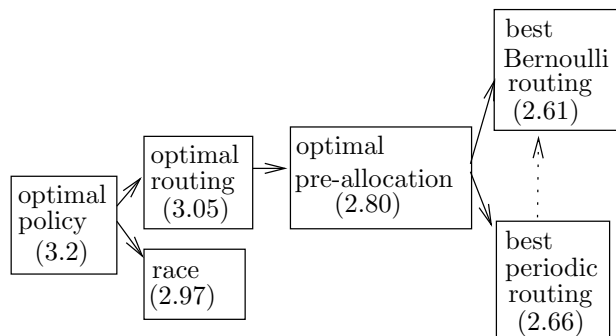
When several transitions of the same cluster belong to r , the decision r induces a race between these transitions.

Example 6. *Let us continue with the Petri net of Figure 2. The optimal non-anticipative policy is as follows:*

- Allocate the token to transition 2 in marking $\{b, e\}$.
- Play race between transitions 2 and 3 in $\{b, d\}$.

The corresponding throughput is approximately 3.20.

The following diagram shows how the different policies compare in terms of throughput. An arc means “provides a better throughput than”. No arcs means that no comparison is possible (several examples of free-choice nets where one or the other policy provides a better throughput have been constructed). The dashed arrow is a conjecture based on numerical evidence. The respective throughputs obtained for the example of Figure 2 are given in parenthesis.



The net in Figure 6 shows that the race policy cannot be compared with the optimal routing: in the previous examples, it has already been shown that the optimal routing can lead to a better throughput than the optimal routing. The net is an example that the contrary can also be true.




Figure 6: A net for which the race policy is better than the optimal routing. With the optimal routing, the chosen transition is always the same and is the transition whose expected firing time is the smallest. Then the throughput is the inverse of that expected firing time. For the race policy, the throughput is the inverse of the expectation of the minimum of the firing times of the two transitions.

6 Conclusion

In this article, we proved the existence of the throughput in free-choice nets under several conflict resolution policies and compared those policies. On the other hand, we exhibit and showed how to compute some optimal policies, according to the information that is available. There, we showed on an example that the optimal non-anticipative policy consists in controlling the system or letting it evolve according to the race policy, depending on the state of the network.

The free-choice assumption was necessary to avoid deadlocks with routing policies, but this assumption can be relaxed if optimal policies are considered.

References

- [1] M. Ajmone-Marsan, G. Balbo, S. Donatelli, G. Franceschinis, and G. Conte. *Modelling with generalized stochastic Petri nets*. Wiley Series in Parallel Computing. Wiley, New-York, 1995.
- [2] E. Altman, B. Gaujal, and A. Hordijk. *Discrete-Event Control of Stochastic Networks: Multimodularity and Regularity*. Number 1829 in Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2003.
- [3] 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.
- [4] F. Baccelli and Foss S. On the saturation rule for the stability of queues. *J. Appl. Probab.*, 32(2):494–507, 1995.
- [5] D. P. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, 1995.
- [6] A. Bouillard, B. Gaujal, and J. Mairesse. Extremal throughputs in free-choice nets. In G. Ciardo and P. Darondeau, editors, *26th International Conference On Application and Theory of Petri Nets*, LNCS 3536. Springer-Verlag, 2005. To appear in *Discrete Event Dyn. Syst.*
- [7] A. Bouillard, B. Gaujal, and J. Mairesse. Throughputs in stochastic routed free-choice nets, existence, computations and optimizations. Technical report, LIP, ENS Lyon, 2005.
- [8] T. Bousch and J. Mairesse. Asymptotic height optimization for topical IFS, Tetris heaps, and the finiteness conjecture. *J. Am. Math. Soc.*, 15(1):77–111, 2002.
- [9] G. Cohen, S. Gaubert, and J.P. Quadrat. Algebraic system analysis of timed Petri nets. In J. Gunawardena, editor, *Idempotency*. Cambridge University Press, 1998.
- [10] B. Gaujal and A. Giua. Timed continuous Petri nets and optimization via linear programming. Technical Report RR-5483, INRIA, 2002.
- [11] B. Gaujal, S. Haar, and J. Mairesse. Blocking a transition in a free choice net and what it tells about its throughput. *J. Comput. Syst. Sci.*, 66(3):515–548, 2003.

-
- [12] T. Mulder and A. Storjohann. On lattice reduction for polynomial matrices. *Journal of symbolic computation*, 35:277–401, 2003.
 - [13] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*. Wiley, 2002.
 - [14] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons Inc., New York, 1994.
 - [15] L. Recalde and M. Silva. Petri net fluidification revisited: Semantics and steady state. *European Journal of Automation APII-JESA*, 35(4):435–449, 2001.



Unité de recherche INRIA Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399