

Throughput in stochastic free-choice nets under various policies

Anne Bouillard*, Bruno Gaujal†, and Jean Mairesse‡

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.

I. 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 [7] and [11]. In both cases, it has been proved that the throughput is simply the solution of a linear program ([8], [11]). The discrete case is more involved. The *deterministic discrete* free-choice case has been studied in [4] 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 [9] 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 , $k \in \{race, Ber, per\}$, 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 [6].

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.

Due to lack of space, most proofs are not reported in this paper. They are given in the long version of the paper, available as a technical report [5].

II. 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

*LIP (UMR CNRS, ENS Lyon, INRIA, Université Claude Bernard Lyon 1), École Normale Supérieure de Lyon, 46, allée d'Italie - F-69364 Lyon Cedex 07 - France. Email: anne.bouillard@ens-lyon.fr

† Lab. ID-MAG, INRIA-CNRS-UJF-INPG, 51, av. J. Kuntzman - F-38330 Montbonnot - France. Email: Bruno.Gaujal@imag.fr

‡LIAFA, CNRS-Université Paris 7, Case 7014, 2 place Jussieu - F-75251 Paris cedex 05 - France. Email: Jean.Mairesse@lafa.jussieu.fr

$M_0 \in \mathbb{N}^{\mathcal{P}}$. 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 $\varphi_a, a \in \mathcal{T}$, 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 [9] 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.

III. 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: (sketch). The result was proved in [9] for the Bernoulli routing. 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.

The case of periodic routings can be proved by adapting the proof of the Bernoulli case. ■

A. 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:

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

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.

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 [9, Prop. 5.1] in which Bernoulli-routed nets are considered.

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.

B. 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 $C^k := \sum_{a \in \mathcal{T}} \lambda_a^k$, $k \in \{race, Ber, per\}$.

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

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

IV. 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 $\mu_a \in (0, \infty)$.




Fig. 1. Example of a live and bounded free-choice Petri net.

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

$$\forall M_1 \in \mathcal{R}(M_0), \quad Q_{M_1, M_1 \cdot a} = \mu_a \text{ if } a \in T_{M_1}$$

$$Q_{M_1, M_1} = - \sum_{a \in T_{M_1}} Q_{M_1, M_1 \cdot a}$$

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 1. 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 $C^{race} = \sum_i -\pi_i \cdot Q_{ii} = 1120/377 \approx 2.97$.

B. 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 $\mathcal{E} = \{(M, T) \mid M \in \mathcal{R}(M_0), T \in \mathcal{T}(M)\}$, *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 1. 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.

C. Periodic routings

We assume that place p_i has a periodic routing policy with period d_i . The behavior of the net can be modeled 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 1. The results are displayed in Figure 2.




Fig. 2. 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 [6].

Numerical evidence also suggests that the best periodic routing is better than the best Bernoulli routing, see Figure 3.

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.

The shape of the curve for the best periodic routing in Figure 3 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 [6] for a closed free-choice Petri net with deterministic timings.




Fig. 3. Comparison of the Bernoulli routing with the best periodic routing when the proportion of tokens sent to transition 2 varies.

V. OPTIMAL ROUTING POLICIES

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.

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 [10]). 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.

A. 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). 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 $\rho \cdot (t, r) = \rho \setminus \{t\} \cup (t \bullet \bullet \cap r(t))$.

Let $J = (J(M, \rho))_{(M, \rho)}$ be the reward vector. Let g be the optimal average reward. The Bellman equation is:

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

Example 3: Consider the example of Figure 1. 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.

B. 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 4: Consider the model of Figure 1. 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 3. 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.

C. 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 5.

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

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

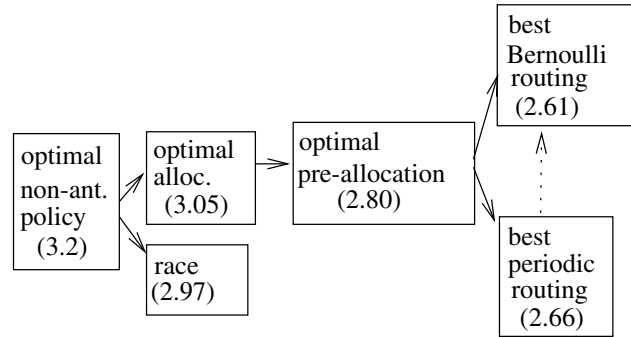
When several transitions of the same cluster belong to r , the decision r induces a race between these transitions.

Example 5: Let us continue with the Petri net of Figure 1. 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 1 are given in parenthesis.



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] 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.*
- [5] A. Bouillard, B. Gaujal, and J. Mairesse. Throughputs in stochastic routed free-choice nets, existence, computations and optimizations. Technical report, LIP, ENS Lyon, 2005.
- [6] 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.
- [7] G. Cohen, S. Gaubert, and J.P. Quadrat. Algebraic system analysis of timed Petri nets. In J. Gunawardena, editor, *Idempotency*. Cambridge University Press, 1998.
- [8] B. Gaujal and A. Giua. Timed continuous Petri nets and optimization via linear programming. Technical Report RR-5483, INRIA, 2002.
- [9] 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.
- [10] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons Inc., New York, 1994.
- [11] L. Recalde and M. Silva. Petri net fluidification revisited: Semantics and steady state. *European Journal of Automation APII-JESA*, 35(4):435–449, 2001.