Equilibrium Analysis of Markov Regenerative Processes

András Horváth[1], Marco Paolieri[2], and Enrico Vicario[3]

[1] University of Turin, Department of Computer Science, Italy
horvath@di.unito.it
[2] University of Southern California, Department of Computer Science, USA
paolieri@usc.edu
[3] University of Florence, Department of Information Engineering, Italy
enrico.vicario@unifi.it

Abstract. We present a solution to compute equilibrium probability density functions (PDFs) for the continuous component of the state in Markov regenerative processes, a class of non-Markovian processes. Equilibrium PDFs are derived as closed-form analytical expressions by applying the Key Renewal Theorem to stochastic state classes computed between regenerations. The solution, evaluated experimentally through the development of an analysis tool, provides the basis to analyze system properties from the equilibrium.

1 Introduction

Stochastic models of discrete-event systems provide a powerful tool for the evaluation of system designs: concurrent activities with stochastic duration can represent service times, arrivals, server breakdowns, or repair actions, for example. Several high-level modeling formalisms are available, including queueing networks [11], stochastic Petri nets [9], stochastic process algebras [5]. Performance and reliability metrics can be evaluated in these models from the transient or steady-state probabilities of their underlying stochastic processes [4].

Steady-state probabilities are computed for the *discrete* component of the system state, such as the number of customers in each queue, or the number of failed servers. These probabilities provide a complete characterization at equilibrium for continuous-time Markov chains (CTMCs), but not for non-Markovian processes, where future evolution also depends on the distribution of the *continuous* component of the state (e.g., remaining time to a failure or service completion).

In fact, CTMCs are memoryless [12]: at any time instant, the future evolution of the process is completely characterized by the current discrete state, independently of previous states or sojourn times. For example, in an M/M/1 queue with arrival rate λ and service rate μ, the discrete state is the number n of customers and times to the next arrival and to the next service (if $n > 0$) are always independent, exponential random variables with rates λ and μ, respectively. Non-Markovian processes do not enjoy such properties: the process evolution after time t depends not only on the discrete state, but also on the distribution of timers at time t [12]. Timers with general (i.e., non-exponential) distributions can "accumulate memory" of previous events and sojourn times: in an M/G/1 queue, the distribution of the remaining service time is not known given the number n of customers, but depends on the time since the last service. Multiple general timers that are concurrently enabled become dependent random variables with a joint distribution [16].

In this paper, we propose a solution to evaluate, in addition to the steady-state probability of discrete states, also the equilibrium distribution of the continuous component of each state, the active timers. Our solution is analytical: using the calculus of stochastic state classes [16,10], we compute closed-form expressions for the joint probability density function (PDF) of active timers immediately after each discrete event of a stochastic model; from these, we derive equilibrium PDFs through the construction of a renewal process and application of the Key Renewal Theorem [12]. Our analysis targets Markov regenerative processes (MRGPs): this class of non-Markovian processes satisfies the Markov property at *regeneration points*, which correspond to time instants where the discrete component of the state provides sufficient information to characterize the PDF of active timers, and thus future evolution [12]. Regeneration points occur when all general timers are reset: in the M/G/1 queue example, each service completion corresponds to a regeneration point of the underlying stochastic process, which is an MRGP.

We restrict our analysis to irreducible MRGPs with finite state space and "bounded memory," i.e., such that a new regeneration point is reached w.p.1 after a bounded number of discrete events (in general, MRGPs can produce trajectories without regenerations, if their measure is zero). Semi-Markov processes (SMPs) are a special case of MRGPs where regeneration points are reached after each discrete event. In contrast with MRGPs under enabling restriction [7], we allow multiple general timers to be concurrently enabled [10,13]. We develop an implementation based on the freely available tool ORIS [14]. Our implementation can automatically compute steady-state probabilities and equilibrium PDFs for each stochastic state class of models where timers are deterministic or sampled according to expolynomial PDFs (products of exponentials and polynomials), which include exponential, uniform, triangular, and Erlang distributions.

The equilibrium analysis of MRGPs is an important result, as it characterizes the stochastic process at the time of a random inspection in the long-run. It generalizes the well-known result for the *remaining life* of a renewal process, which has PDF $f_Y(y) = [1 - F_X(y)]/E[X]$ after a random inspection if $F_X(x)$ is the cumulative distribution function (CDF) of inter-event times [12, Eq. (8.40)]. When the inspection represents a catastrophic failure, equilibrium PDFs can be modified to reflect its effects and used in transient analysis to compute survivability metrics, similarly to solutions for CTMCs [15,8,6]. Once the equilibrium PDFs are known for each state, the approach also enables the generation of samples from the equilibrium distribution of the MRGP process without the need to monitor convergence and mixing during a simulation, similarly to *perfect sampling* [1,2,3] methods for DTMCs and CTMCs.

2 Markov Regenerative Processes

2.1 Stochastic Time Petri Nets

We adopt stochastic time Petri nets (STPNs) to specify discrete-event systems governed by stochastic timers. We refer to Appendix A for a complete definition of STPNs and present only their essential elements. An STPN includes a set P of *places* (graphically drawn as circles) and a set T of *transitions* (drawn as vertical bars): transitions represent concurrent activities that move *tokens* between places.

State. The state $s = (m, \vec{\tau})$ of an STPN includes two components: (1) a *marking* $m\colon P \to \mathbb{N}$ that assigns a *token count* to each place and controls the enabling of transitions, and (2) a *time-to-fire* vector $\vec{\tau} = (\tau_1, \ldots, \tau_n)$ that specifies the remaining time $\tau_i \in \mathbb{R}_{\geqslant 0}$ to the *firing* of each transition enabled by m (given a total order on T).

State Update. A transition is enabled by marking m if all input places (connected with incoming arcs) contain at least one token. The enabled transition t^* with minimum remaining time in $\vec{\tau}$ fires and produces a new state $s' = (m', \vec{\tau}')$ where: m' is obtained from m by removing one token from each input place of t^* and adding one token to each output place of t^* (connected with outgoing arcs). Transitions enabled before and after each step, but distinct from t^*, are called *persistent*: after the firing, their remaining times to the fire are reduced by that of t^*, i.e., ${\tau_i}' = \tau_i - \tau_{t^*}$. Other transitions enabled by m' are called *newly-enabled*: their remaining times ${\tau_i}'$ are sampled independently according to CDFs $F_t(x)$ specified by the STPN for each transition t.

We assume that each CDF F_t admits the representation $F_t(x) = \int_0^x f_t(u)\, du$. We represent the PDF of a transition t with deterministic duration $\overline{x}$ using the Dirac delta function $f_t(x) = \delta(x - \overline{x})$. As usual in stochastic Petri nets, the PDF family of a transition is represented graphically using white rectangles for exponential transitions, gray rectangles for deterministic ones, black rectangles for other distributions.

Example 1 (Parallel Producer-Consumer). Fig. 1a presents the STPN model of two producers working in parallel to produce parts that are consumed together by a single consumer; consumption begins only when both parts are available, and production of new parts starts when the previous ones have been consumed. Tokens in places $\texttt{pIn}_1$ and $\texttt{pIn}_2$ activate the two producers represented by transitions $\{\texttt{prod}_{11}, \texttt{prod}_{12}\}$ and $\texttt{prod}_2$, respectively. The first producer uses two processing units to increase performance: the first among $\texttt{prod}_{11}$ and $\texttt{prod}_{12}$ to complete (respectively, with time to fire PDF $f(x) = 2x - 2$ on $[1, 2]$ and uniform on $[1, 2]$) ends the production of the first part. The second producer is modeled by transition $\texttt{prod}_2$ with PDF $f(x) = 1 - x/2$ on $[0, 2]$. The consumer, modeled by transition `consume` with deterministic firing time equal to 1, is enabled when both tokens are moved to places $\texttt{pOut}_1$ and $\texttt{pOut}_2$; after its completion, tokens are moved back to places $\texttt{pIn}_1$ and $\texttt{pIn}_2$, and production restarts. Fig. 1b presents the *state class graph* [16] for this model, where edges represent possible transitions firings and each node S_i represents the marking m_i and the set D_i of possible values for the time-to-fire vector $\vec{\tau}$ immediately after a firing.

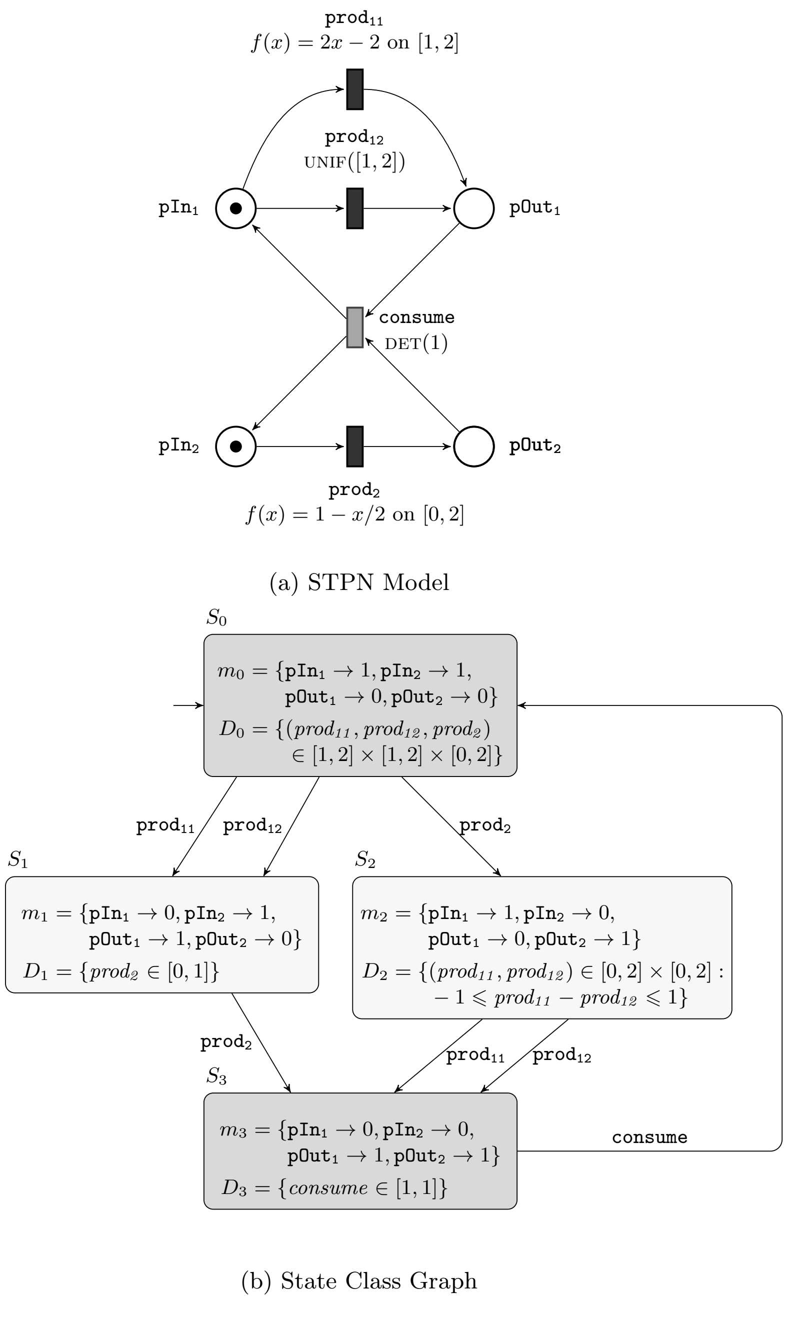

Fig. 1: Parallel Producer-Consumer Example

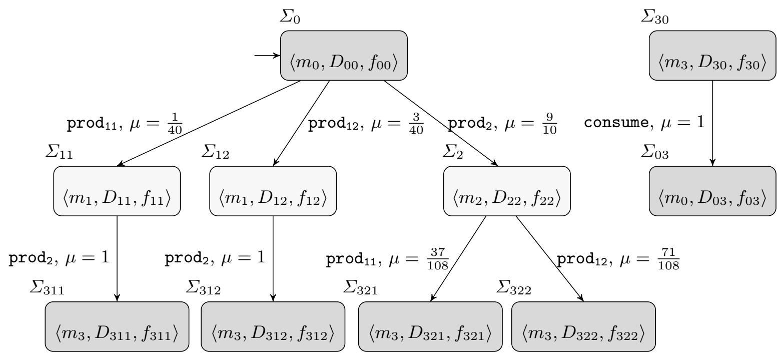

Fig. 2: Trees of Stochastic State Classes from Σ_0 and Σ_{00}

2.2 Markov Regenerative Processes

Given an initial marking m_0 and a joint PDF $f_0(\vec{x})$ of initial times to fire $\vec{\tau}_0$, each execution of the STPN produces a sequence of state changes $s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} s_2 \xrightarrow{t_3} \cdots$ where $s_0 = (m_0, \vec{\tau}_0)$ is the initial state, $t_i \in T$ is the ith fired transition, and $s_i = (m_i, \vec{\tau}_i)$ is the state reached after the firing of t_i. Firing of a transition is a *regeneration point* if all general (i.e., non-exponential) transitions that are enabled after the firing are resampled (newly-enabled) which means that the marking reached due to the firing provides sufficient information to reconstruct the PDF of the time-to-fire vector (which is simply the product of PDFs of enabled transitions). On the other hand, firing of a transition is not a regeneration point if one or more general transitions remain enabled. The *marking process* $\{\, Z(t),\ t \geqslant 0 \,\}$ records the marking of the STPN as it evolves over time. It is a continuous-time process with a countable state space, the set of markings $M \subseteq \mathbb{N}^P$ (for a formal definition, see [9, Sect. 3.1]). The family of the marking process depends on the type of time-to-fire distributions and on the overlap of time intervals during which general transitions are enabled [4]. We focus on the class of MRGPs that allows multiple general transitions to be enabled at the same time but requires that, in a bounded number of transition firings, the model reaches a regeneration point; we denote the set of markings reached at regeneration points as $R \subseteq M$.

Example 2 (Parallel Producer-Consumer). Regeneration points for the marking process of the STPN in Fig. 1a are highlighted in the state class graph of Fig. 1b by darker backgrounds; they correspond to the firings that lead to marking m_3 (i.e., when production ends) or to marking m_0 (i.e., when `consume` fires and production restarts), i.e., $R = \{m_3, m_0\}$.

MRGPs provide a good trade-off between modeling power and complexity of the analysis: concurrent general timers can persist to discrete events, while transient and steady-state probabilities can be computed numerically [10,13].

$$
\begin{aligned}
&f_{00}(age, prod_{11}, prod_{12}, prod_2) = (2prod_{11} - prod_{11}\, prod_2 + prod_2 - 2)\ \delta(age)\\
&D_{00} = \{(age, prod_{11}, prod_{12}, prod_2) \in [0,0] \times [1,2] \times [1,2] \times [0,2]\}\\
&f_{11}(age, prod_2) = -200age^2 - 320age + 40age^2\, prod_2 +\\
&\qquad 120age\, prod_2 - 40age^3 - 160 + 80prod_2\\
&D_{11} = \{(age, prod_2) \in [-2,-1] \times [0,1] : -2 \leqslant age - prod_2 \leqslant -1\}\\
&f_{12}(age, prod_2) = -\frac{80}{3}age^2 - \frac{80}{3}age + \frac{20}{3}age^2 prod_2 + \frac{40}{3}age\, prod_2 - \frac{20}{3}age^3\\
&D_{12} = \{(age, prod_2) \in [-2,-1] \times [0,1] : -2 \leqslant age - prod_2 \leqslant -1\}\\
&f_{311}(age, consume) = \left(\frac{200}{3} + \frac{580}{3}age + 200age^2 + \frac{260}{3}age^3 + \frac{40}{3}age^4\right)\delta(consume - 1)\\
&D_{311} = \{(age, consume) \in [-2,-1] \times [1,1]\}\\
&f_{312}(age, consume) = \left(-\frac{80}{9} - \frac{40}{9}age + \frac{40}{3}age^2 + \frac{100}{9}age^3 + \frac{20}{9}age^4\right)\delta(consume - 1)\\
&D_{312} = \{(age, consume) \in [-2,-1] \times [1,1]\}
\end{aligned}
$$

Fig. 3: PDFs and supports for paths $\Sigma_0 \to \Sigma_{11} \to \Sigma_{311}$ and $\Sigma_0 \to \Sigma_{12} \to \Sigma_{312}$

Transient probabilities $P_{ij}(t) := P(Z(t) = j \mid X_0 = i)$ from all initial regenerations $i \in R$ and for all $j \in M$, $t \geqslant 0$ can be computed from the system of Markov renewal equations [12]

$$\mathbf{P}(t) = \mathbf{L}(t) + \int_0^t d\mathbf{G}(u)\, \mathbf{P}(t-u) \tag{1}$$

where

$$G_{ik}(t) := P(X_1 = k,\, T_1 \leqslant t \mid X_0 = i)$$

for $i, k \in R$ is the *global kernel* of the MRGP, specifying the joint distribution of the next regeneration X_1 and regeneration point T_1 given that the last regeneration was $X_0 = i$ at time 0, while

$$L_{ij}(t) := P(Z(t) = j, T_1 > t \mid X_0 = i)$$

is the *local kernel* of the MRGP, defined as the probability that, given the initial regeneration $i \in R$ at time 0, no further regeneration has been reached and the marking is $j \in M$ at time t. Informally, the global kernel describes the process of the regeneration points while the local kernel provides the necessary information between two consecutive regeneration points. The system of Eq. (1) is a set of Volterra integral equations that can be solved numerically in the time domain.

The global and local kernels also provide the steady-state probabilities p_j of each marking $j \in M$. If $\vec{\pi}$ is the vector of steady-state probabilities of the discrete-time Markov chain (DTMC) embedded at regeneration points, i.e., $\sum_{k \in R} \pi_k = 1$ and $\vec{\pi} = \mathbf{G}(\infty)\vec{\pi}$, then

$$p_j = \frac{\sum_{i \in R} \pi_i \alpha_{ij}}{\sum_{i \in R, j' \in M} \pi_i \alpha_{ij'}} \tag{2}$$

where $\alpha_{ij} := \int_0^\infty L_{ij}(t)\, dt$ is the expected time spent in j after a regeneration in $i \in R$ and before the next one [13].

2.3 Analysis with Stochastic State Classes

Although Eqs. (1) and (2) provide an elegant solution to compute transient and steady-state probabilities, a major difficulty lies in the evaluation of the global and local kernels for a given model. One approach is to compute the joint PDF of the time-to-fire vector and firing time after each transition firing until a regeneration. This can be accomplished through the calculus of *stochastic state classes* [16,10].

Definition 1 (Stochastic State Class). *A stochastic state class Σ is a tuple $\langle m, D, f\rangle$ where: $m \in M$ is a marking; f is the PDF (immediately after a firing) of the random vector $\langle \tau_{age}, \vec{\tau}\rangle$ including the time-to-fire vector $\vec{\tau}$ of transitions enabled by m and the age variable τ_{age} accumulating previous sojourn times; $D \subseteq \mathbb{R}^{n+1}$ is the support of f.*

The initial stochastic state class has marking m_0 (the initial marking of the STPN) and PDF $f(x_{age}, \vec{x}) = \delta(x_{age}) f_0(\vec{x})$, where f_0 is the PDF of the initial time-to-fire vector $\vec{\tau}_0$ of the STPN and δ is the Dirac delta. Given a class $\Sigma = \langle m, D, f\rangle$ and a transition t enabled by m, the calculus of [16] computes (1) the probability μ that t is the transition that fires in Σ, and (2) the *successor class* Σ', which includes the marking and time-to-fire PDF after the firing of t in Σ. In the calculus, the age variable τ_{age} is decreased by the sojourn time, to treat it similarly to persistent times to fire [10]; the time of the last firing is thus given by $-\tau_{age}$.

To analyze MRGPs, regeneration points are detected during the computation of successors and new regeneration states are included in the set R. Each regeneration state $i \in R$ uniquely identifies an initial marking and PDF of the time-to-fire vector, which can be used to construct an initial stochastic state class. By computing trees of stochastic state classes from each $i \in R$ to other regeneration points, the MRGP is encoded as a set of *trees* of stochastic state classes.

If $\textsc{Inner}(i)$ and $\textsc{Leaves}(i)$ are, respectively, the stochastic state classes of inner nodes and leaf nodes in the transient tree enumerated from regeneration $i \in R$, then

$$L_{ij}(t) = \sum_{\substack{\Sigma \in \textsc{Inner}(i) \text{ s.t.} \\ \Sigma \text{ has marking } j}} p_{in}(\Sigma, t) \quad \text{and} \quad G_{ik}(t) = \sum_{\substack{\Sigma \in \textsc{Leaves}(i) \text{ s.t.} \\ \Sigma \text{ has regeneration } k}} p_{reach}(\Sigma, t)$$

for all $i, k \in R$, $j \in M$, and $t \geqslant 0$, where for a class $\Sigma = \langle m, D, f\rangle$ reached through firings with probability $\rho(\Sigma)$,

$$p_{reach}(\Sigma, t) = \rho(\Sigma) \int_{\{(x_{age}, \vec{x}) \in D \mid -x_{age} \leqslant t\}} f(x_{age}, \vec{x})\, dx_{age}\, d\vec{x}$$

is the probability that Σ is reached from i within time t and

$$p_{in}(\Sigma, t) = \rho(\Sigma) \int_{\{(x_{age}, \vec{x}) \in D \mid -x_{age} \leqslant t \text{ and } x_k - x_{age} > t\ \forall k\}} f(x_{age}, \vec{x})\, dx_{age}\, d\vec{x}$$

is the probability that all and only the transitions leading from i to Σ have fired by time t [10]. When timers are deterministic or with expolynomial PDFs, the integrals in the two equations above (and similar integrals in the rest of the paper) can be evaluated exactly in ORIS by symbolic integration over zones.

Example 3 (Parallel Producer-Consumer). Fig. 2 presents the trees of stochastic state classes for the example of Fig. 1a. Regenerations are identified by markings $R = \{m_0, m_3\}$: $\textsc{Inner}(m_0) = \{\Sigma_0, \Sigma_{11}, \Sigma_{12}, \Sigma_2\}$, $\textsc{Inner}(m_3) = \{\Sigma_{30}\}$, while $\textsc{Leaves}(m_0) = \{\Sigma_{311}, \Sigma_{312}, \Sigma_{321}, \Sigma_{322}\}$, and $\textsc{Leaves}(m_3) = \{\Sigma_{03}\}$. Edges are labeled with firing probabilities, e.g., the probability $\rho(\Sigma_{321})$ of firing $\texttt{prod}_2$ and then $\texttt{prod}_{11}$ is $\frac{9}{10} \cdot \frac{37}{108}$.

3 Equilibrium Analysis

3.1 Steady-state Probabilities

To compute the probability of observing each stochastic state class in steady state, we follow the strategy of Eq. (2) but consider each class (instead of each marking) as a distinct state j. First, we compute the limit of the global kernel $\mathbf{G}(t)$ as $t \to \infty$. Each entry $G_{ik}(\infty)$ for $i, k \in R$ can be obtained as the product of firing probabilities from regeneration i to all classes in $\textsc{Leaves}(i)$ that reach regeneration $k \in R$:

$$G_{ik}(\infty) := P(X_1 = k \mid X_0 = i) = \sum_{\substack{\Sigma \in \textsc{Leaves}(i) \text{ s.t.} \\ \Sigma \text{ has regeneration } k}} \rho(\Sigma)$$

where $\rho(\Sigma)$ is the product of firing probabilities μ for the sequence of firings that leads from $i \in R$ to Σ.

Next, for each regeneration $i \in R$ and class $j \in \textsc{Inner}(i)$, we compute the expected time α_j spent in j after reaching i and before the next regeneration.

Lemma 1. *Let $j = \langle m, D, f \rangle$ be an inner node in the tree enumerated from regeneration $i \in R$, i.e., $j \in \textsc{Inner}(i)$. Then, the expected sojourn time of the MRGP in j before the next regeneration is*

$$\alpha_j = \rho(j) \sum_{t \in E(j)} \mu^{(t)} \int_{D^{(t)}} x_{k_t} \, f^{(t)}(x_{age}, \vec{x}) \, dx_{age} \, d\vec{x}, \tag{3}$$

where: $\rho(j)$ is the product of firing probabilities of transitions that lead from regeneration i to class j; $E(j) \subseteq T$ is the set of transitions enabled in j; $\mu^{(t)}$ is the probability that $t \in E(j)$ fires in j; $D^{(t)} = \{(x_{age}, \vec{x}) \in D \mid x_k \geqslant x_{k_t} \ \forall k\}$ is the subset of the support D where τ_{k_t}, the time to fire of t, is minimum (k_t is the index of t in $\vec{\tau}$); and

$$f^{(t)}(x_{age}, \vec{x}) := f(x_{age}, \vec{x}) \left(\int_{D^{(t)}} f(x_{age}, \vec{x}) \, dx_{age} \, d\vec{x} \right)^{-1} \tag{4}$$

is the PDF of $\langle \tau_{age}, \vec{\tau} \rangle$ conditioned on $\{\tau_{k_t}$ is minimum$\}$.

Proof. Eq. (3) follows from the definition of stochastic state class and from the law of total expectation. The events of the firing in j of transitions in $E(j)$ are mutually exclusive and exhaustive, and thus, if S_j is the sojourn time in j, $\alpha_j := \rho(j)\, E[S_j] = \rho(j) \sum_{t \in E(j)} \mu^{(t)}\, E[S_j \mid t \text{ fires in } j]$.

Similarly to Eq. (2), the steady-state probability of class $j \in \textsc{Inner}(i)$ is

$$p_j = \frac{\pi_i \alpha_j}{\sum_{i' \in R, j' \in \textsc{Inner}(i')} \pi_{i'} \alpha_{j'}} \tag{5}$$

where $\vec{\pi}$ is such that $\sum_{k \in R} \pi_k = 1$ and $\vec{\pi} = \mathbf{G}(\infty)\vec{\pi}$.

Example 4 (Parallel Producer-Consumer). The MRGP of Fig. 2 has a simple DTMC embedded at regeneration points: the process alternates between regenerations $R = \{m_0, m_3\}$, and thus $\mathbf{G}(\infty) = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $\vec{\pi} = (0.5, 0.5)$. The sojourn times α_j of inner nodes are: $(\frac{1}{40}\frac{4}{3} + \frac{3}{40}\frac{11}{9} + \frac{9}{10}\frac{31}{54}) = \frac{77}{120}$ for $j = \Sigma_0$, $\frac{1}{40}(\frac{2}{9}) = \frac{1}{180}$ for $j = \Sigma_{11}$, $\frac{3}{40}(\frac{7}{27}) = \frac{7}{360}$ for $j = \Sigma_{12}$, $\frac{9}{10}(\frac{37}{108}\frac{34}{37} + \frac{71}{108}\frac{59}{71}) = \frac{31}{40}$ for $j = \Sigma_2$, 1 for $j = \Sigma_{30}$. These give steady-state probabilities p_j equal to $\frac{77}{293}$, $\frac{2}{879}$, $\frac{7}{879}$, $\frac{93}{293}$, $\frac{120}{293}$ for $j = \Sigma_0, \Sigma_{11}, \Sigma_{12}, \Sigma_2, \Sigma_{30}$, respectively.

3.2 Equilibrium PDFs

Marginal PDF of $\vec{\tau}$ in Σ given the firing of t_1. Without loss of generality, we assume that n transitions $t_1, \dots, t_n$ are enabled in Σ and consider the case where t_1 is the one that fires. Conditioned on this event, the PDF of $\vec{\tau}$ in Σ is

$$f^{(t_1)}(\vec{x}) = \int_{D^{(t_1)}} f(x_{age}, \vec{x})\, dx_{age} \left(\int_{D^{(t_1)}} f(x_{age}, \vec{x})\, dx_{age}\, d\vec{x} \right)^{-1} \tag{6}$$

where $D^{(t_1)} = \{(x_{age}, \vec{x}) \in D \mid x_k \geqslant x_1\ \forall k\}$ is the subset of the support D where τ_1 is minimum. Eq. (6) follows by restricting the support to the subset $D^{(t_1)}$, normalizing the PDF $f(x_{age}, \vec{x})$, and then obtaining the marginal PDF of $\vec{\tau}$ by integrating over all possible values for τ_{age} in $D^{(t_1)}$.

Stochastic process $\vec{r}(t)$ of times to fire $\vec{\tau}$ across renewals. Successive visits to the stochastic state class $\Sigma = \langle m, D, f \rangle$ observe copies of $\langle \tau_{age}, \vec{\tau} \rangle$ that are independent and identically distributed (i.i.d.) according to the PDF f, since the MRGP encounters a regeneration point between visits and then performs the same sequence of transition firings. Similarly, visits to Σ that end with the firing of t_1 also observe the same PDF $f^{(t_1)}$ of $\vec{\tau}$ derived in Eq. (6), and their sojourn times are i.i.d. random variables.

We focus our attention on the time intervals of these i.i.d. sojourn times: as time advances, we move from a sojourn in Σ to the next one, always under the hypothesis that t_1 is the transition that fires in Σ. We construct a renewal process $\{N(t), t \geqslant 0\}$ where times between events (i.e., interarrival times) are distributed as a sojourn in Σ that ends with the firing of t_1. We denote interarrival times as $S_1, S_2, \dots$ and renewal times as $T_k = \sum_{i=1}^{k} S_i$ for $k \geqslant 0$; $N(t) = \max\{k \mid T_k \leqslant t\}$

is the number of sojourns completed by time t and $N(t) = k \Leftrightarrow T_k \leqslant t < T_{k+1}$. The interarrival PDF of this renewal process is given by the marginal PDF of τ_1 given that it is minimum in $\vec{\tau}$, which is

$$g(x_1) = \int_{D^{(t_1)}} f^{(t_1)}(x_1, x_2, \ldots, x_n)\, dx_2 \cdots dx_n \tag{7}$$

where $D^{(t_1)}$ is the support of $f^{(t_1)}$. Eq. (7) integrates over all possible values of $\tau_2, \ldots, \tau_n$ to obtain the marginal PDF of τ_1 when $\tau_1 \leqslant \tau_i$ for all $i = 2, \ldots, n$.

As $N(t)$ evolves across each renewal $T_0, T_1, T_2, \ldots$, a new time-to-fire vector $\vec{\tau}^{(i)}, i = 0, 1, 2, \ldots$ is sampled independently at each T_i according to the same PDF $f^{(t_1)}$. Our goal is to study the evolution of these time-to-fire random vectors over time, subject to the fact that renewal times T_i are also random. We denote by $\{\, \vec{r}(t),\ t \geqslant 0 \,\}$ the n-dimensional stochastic process describing, for each $t \geqslant 0$, the current value of the time-to-fire vector, i.e., $\vec{r}(t) := \vec{\tau}^{(N(t))} - (t - T_{N(t)})$, and denote its PDF at all $t \geqslant 0$ by $h(t, \vec{x})$, i.e.,

$$P(r_1(t) \leqslant x_1, \ldots, r_n(t) \leqslant x_n) := \int_{-\infty}^{x_1} \cdots \int_{-\infty}^{x_n} h(t, \vec{x})\, dx_1 \cdots dx_n .$$

Equilibrium PDF of $\vec{r}(t)$. Our goal is to compute the equilibrium PDF of $\vec{r}(t)$, i.e., the function $\hat{f}^{(t_1)}(\vec{x}) = \lim_{t\to\infty} h(t, \vec{x})$, which gives the PDF of the times to fire in Σ at equilibrium (given that sojourns end with the firing of t_1). First, we provide the following result, which highlights the fundamental relation between $h(t, \vec{x})$, the object of our analysis, and the PDF $f^{(t_1)}(\vec{x})$, which can be readily computed using Eq. (6).

Lemma 2 (Renewal Equation for h). *If $h(t, \vec{x})$ is the PDF of $\vec{r}(t)$ for each $t \geqslant 0$, $f^{(t_1)}(\vec{x})$ is the PDF of $\vec{\tau}$ at each renewal, and $g(x)$ is the PDF of τ_1 (the interarrival time of the renewal process), the following renewal equation holds:*

$$h(t, \vec{x}) = f^{(t_1)}(\vec{x} + t) + \int_0^t h(t - u, \vec{x})\, g(u)\, du\, . \tag{8}$$

Proof. Eq. (8) can be derived through a renewal argument: for the first renewal time S_1 we have that either $S_1 > t$ or $S_1 \leqslant t$.

If $S_1 > t$, then the first renewal has not occurred, so that $N(t) = 0$ and $\vec{r}(t) = \vec{\tau}^{(0)} - t$. The PDF of $\vec{r}(t)$ at time t is then given by $f^{(t_1)}(\vec{x} + t)/P(\tau_1 > t)$, i.e., the PDF $f^{(t_1)}$ used to sample $\vec{\tau}^{(0)}$ but conditioned on the event $\{\tau_1 > t\}$ and where each component is shifted by time t (we denote by $\vec{x} + t$ the vector $(x_1 + t, \ldots, x_n + t)$). Then, we have that $h(t, \vec{x} \mid S_1 > t)\, P(S_1 > t) = f^{(t_1)}(\vec{x} + t)$, since $S_1 := \tau_1$.

If $S_1 \leqslant t$, the process $\vec{r}(t)$ "probabilistically restarts" after S_1, when a new time-to-fire vector $\vec{\tau}^{(1)}$ is sampled. Formally, if $S_1 = u$, at least one renewal is encountered by time t, $N(t) = N(t - u) + 1$, $T_{N(t-u)+1} = T_{N(t-u)} + u$, and thus

$$\begin{aligned} \vec{r}(t) &= \vec{\tau}^{(N(t-u)+1)} - (t - T_{N(t-u)+1}) \\ &= \vec{\tau}^{(N(t-u)+1)} - [(t - u) - T_{N(t-u)}] \end{aligned}$$

for $u \leqslant t$. Given that time-to-fire vectors $\vec{\tau}^{(N(t-u)+1)}$ and $\vec{\tau}^{(N(t-u))}$ have the same PDF $f^{(t_1)}$, it holds that $h(t,\vec{x}) = h(t-u,\vec{x})$ for $u \leqslant t$. By conditioning on all the possible values of $S_1 = u$ and decreasing t accordingly, we have

$$h(t,\vec{x} \mid S_1 \leqslant t)\, P(S_1 \leqslant t) = \int_0^t h(t-u,\vec{x})\, g(u)\, du\,.$$

By putting together the two cases, we obtain Eq. (8).

Lemma 2 establishes a connection between h and $f^{(t_1)}$ and also reveals the recursive structure of h across renewals. This kind of renewal-type equation is well-known for renewal processes and provides a strategy to compute h at the equilibrium through the following result [12, Theorem 8.17].

Theorem 1 (Key Renewal Theorem). *Let $g(x)$ be the PDF of the interarrival time, and let h be a solution to the renewal-type equation $h(t) = d(t) + \int_0^t h(t-u)g(u)\, du$. Then, if d is the difference of two non-negative bounded monotone functions and $\int_0^\infty |d(u)|\, du < \infty$,*

$$\lim_{t\to\infty} h(t) = \frac{1}{E[S]} \int_0^\infty d(u)\, du$$

where $E[S] = \int_0^\infty u\, g(u)\, du$ is the mean interarrival time.

Theorem 1 applies to Eq. (8) with $d(t) = f^{(t_1)}(\vec{x}+t)$. Moreover, when the PDFs f_t used to sample newly-enabled transitions are piecewise expolynomials (products of exponentials and polynomials), the joint PDF f of timers, and thus $f^{(t_1)}$, is also piecewise continuous and with bounded variation [16].

By combining Lemma 2 and Theorem 1, we obtain the equilibrium PDF $\hat{f}^{(t_1)}$ of $\vec{\tau}$ in Σ when t_1 is the transition that fires at the end of each sojourn:

$$\hat{f}^{(t_1)}(\vec{x}) := \lim_{t\to\infty} h(t,\vec{x}) = \frac{1}{E[S^{(t_1)}]} \int_0^\infty f^{(t_1)}(\vec{x}+u)\, du \tag{9}$$

where $E[S^{(t_1)}] = \int_0^\infty u\, g(u)\, du$ is the mean sojourn time in Σ when t_1 fires. The identity of Eq. (9) is a major step for the analysis of the joint PDF of $\vec{\tau}$ at the steady state. Combined with Eq. (6) to obtain $f^{(t_1)}$ from f, and with Eq. (7) to obtain g from $f^{(t_1)}$, it provides a straightforward derivation of the equilibrium PDF under the hypothesis that t_1 is always the transition that fires in Σ.

Equilibrium PDF when multiple transitions can fire. The equilibrium PDF $\hat{f}^{(t_1)}$ of Eq. (9) assumes that, after each visit to Σ, transition t_1 is always the one that fires among $t_1, \ldots, t_n$. The following theorem removes this hypothesis.

Theorem 2 (Equilibrium PDF). *Let $\Sigma = \langle m, D, f\rangle$ be a stochastic state class where transitions $t_1, \ldots, t_n$ can fire with probabilities $\mu^{(t_1)}, \ldots, \mu^{(t_n)}$, respectively. Then, the equilibrium PDF of $\vec{\tau} = (\tau_1, \ldots, \tau_n)$ is given by*

$$\hat{f}(\vec{x}) = \frac{1}{E[S]} \sum_{i=1}^{n} \mu^{(t_i)} \int_0^\infty f^{(t_i)}(\vec{x}+u)\, du \tag{10}$$

where $E[S]$ *is the expected sojourn time in* Σ *and, for all* $i = 1, \dots, n$, $f^{(t_i)}$ *is the PDF of* $\vec{\tau}$ *conditioned on the firing of* t_i *according to Eq.* (6).

Proof. We focus only on sojourns in class Σ and ignore the rest of the time line. The probability that a sojourn ends with the firing of t_i is $\mu^{(t_i)}$ for $i = 1, \dots, n$, with $\sum_{i=1}^{n} \mu^{(t_i)} = 1$; conditioned on this event, the expected sojourn time in Σ is $E[S^{(t_i)}]$. Then, the steady-state probability of sojourns in Σ that end with the firing of t_i is given by

$$p_i = \frac{\mu^{(t_i)} E[S^{(t_i)}]}{\sum_{j=1}^{n} \mu^{(t_j)} E[S^{(t_j)}]}$$

which is the mean fraction of time spent in such sojourns. Since $\hat{f}^{(t_i)}$ is the equilibrium PDF when sojourns end with t_i,

$$\begin{aligned}
\hat{f}(\vec{x}) &= \sum_{i=1}^{n} p_i \, \hat{f}^{(t_i)}(\vec{x}) = \sum_{i=1}^{n} \left(\frac{\mu^{(t_i)} E[S^{(t_i)}]}{\sum_{j=1}^{n} \mu^{(t_j)} E[S^{(t_j)}]} \right) \hat{f}^{(t_i)}(\vec{x}) \\
&= \frac{1}{\sum_{j=1}^{n} \mu^{(t_j)} E[S^{(t_j)}]} \sum_{i=1}^{n} \mu^{(t_i)} \int_0^\infty f^{(t_i)}(\vec{x} + u) \, du
\end{aligned}$$

which, since $\sum_{j=1}^{n} \mu^{(t_j)} E[S^{(t_j)}] = E[S]$, gives Eq. (10).

4 Experimental Evaluation

Steady-state probabilities and equilibrium PDFs represent the equilibrium distribution of the MRGP. When used as an initial distribution for transient analysis, this distribution must result in constant transient probabilities that are equal to the steady-state ones. In this section, we describe how to perform transient analysis from this distribution and validate the correctness of the approach.

In Section 3.1, we derived the steady-state probability p_c of each each class $c \in \cup_{i \in R} \textsc{Inner}(i)$. Given that the MRGP is in class $c = \langle m, D, f \rangle$, the marking is equal to m and the time-to-fire vector $\vec{\tau}$ has equilibrium PDF given by $\hat{f}(\vec{x})$, which is computed from f according to Eq. (10) of Theorem 2. To compute transient probabilities from the equilibrium, we modify the approach of Eq. (1) as follows.

First, for each inner node $c = \langle m, D, f \rangle, c \in \cup_{i \in R} \textsc{Inner}(i)$, we compute a tree of stochastic state classes until the next regeneration. We construct the initial class $\textsc{Start}(c)$ of this tree using marking m and PDF of $\langle \tau_{age}, \vec{\tau} \rangle$ equal to $g(x_{age}, \vec{x}) = \delta(x_{age}) \hat{f}(\vec{x})$, i.e., $\tau_{age} = 0$ and the time-to-fire vector $\vec{\tau}$ has PDF $\hat{f}(\vec{x})$. For each $c \in \cup_{i \in R} \textsc{Inner}(i)$, we denote the inner nodes and leaves of the tree computed from $\textsc{Start}(c)$ (until the next regeneration) as $\textsc{StartInner}(c)$ and $\textsc{StartLeaves}(c)$, respectively.

Then, we extend the Markov renewal equations of Eq. (1) by introducing an additional regeneration $\hat{r}$ that represents the state of the MRGP at equilibrium. The process starts in $\hat{r}$ at time 0, but never returns to this artificial regeneration:

by construction, the next regeneration belongs to R and, afterward, the MRGP cycles through its original trees of stochastic state classes. To achieve this behavior, we set MRGP kernel entries as follows. Let $\hat{R} = R \cup \{\hat{r}\}$ and set, for $i = \hat{r}$,

$$L_{ij}(t) = \sum_{c \in \cup_{i' \in R} \textsc{Inner}(i')} p_c \left(\sum_{\substack{\Sigma \in \textsc{StartInner}(c) \text{ s.t.} \\ \Sigma \text{ has marking } j}} p_{in}(\Sigma, t) \right) \tag{11}$$

$$G_{ik}(t) = \sum_{c \in \cup_{i' \in R} \textsc{Inner}(i')} p_c \left(\sum_{\substack{\Sigma \in \textsc{StartLeaves}(c) \text{ s.t.} \\ \Sigma \text{ has regeneration } k}} p_{reach}(\Sigma, t) \right) \tag{12}$$

for all $k \in R$, $j \in M$, and $t \geqslant 0$. Since $\hat{r}$ is never reached again, we set $G_{ik}(t) = 0$ for all $i \in \hat{R}$ when $k = \hat{r}$.

Kernel entries in the additional row $\hat{r}$ model a random choice of the initial stochastic state class c according to the discrete distribution given by p_c for $c \in \cup_{i \in R} \textsc{Inner}(i)$; for a given class c, the tree computed from $\textsc{Start}(c)$ is used to characterize the system evolution from the equilibrium in c and until the next regeneration. As in Section 2.3, measures $p_{in}(\Sigma, t)$ and $p_{reach}(\Sigma, t)$ provide the probability that the MRGP is in the stochastic state class Σ at time t, and that it has reached Σ by time t, respectively.

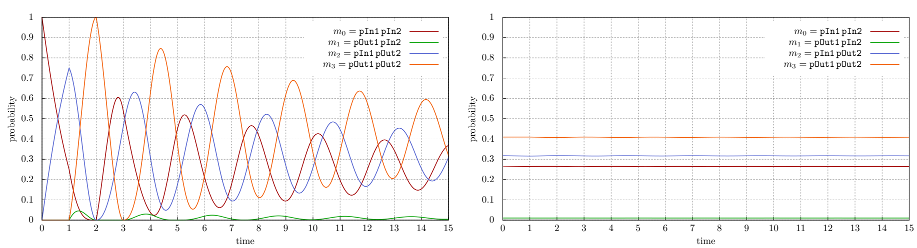

(a) From initial state (b) From equilibrium PDFs

Fig. 4: Transient Analysis of Parallel Producer-Consumer Example

Example 5. We consider the STPN of Fig. 1a and its underlying MRGP of Fig. 2 with markings $M = \{m_0, m_1, m_2, m_3\}$ (defined in Fig. 1b), regenerations $R = \{m_0, m_3\}$, inner nodes $\textsc{Inner}(m_0) = \{\Sigma_0, \Sigma_{11}, \Sigma_{12}, \Sigma_2\}$ and $\textsc{Inner}(m_3) = \{\Sigma_{30}\}$. The steady-state probabilities p_c result in steady-state probabilities of the marking process equal to $\frac{77}{293} \approx 0.263$ for m_0 (steady state probability of Σ_0), $\frac{2}{879} + \frac{7}{879} \approx 0.010$ for m_1 (steady-state probabilities of Σ_{11} and Σ_{12}), $\frac{93}{293} \approx 0.317$ for m_2 (steady-state probability of Σ_2), $\frac{120}{293} \approx 0.410$ for marking m_3 (steady-state probability of Σ_{30}). Fig. 4a illustrates the transient probabilities $P_{ij}(t)$ for $0 \leqslant t \leqslant 15$ of the MRGP for $i = m_0$ (i.e., from the initial regeneration) and for each $j \in M$. Fig. 4b shows the transient probabilities $P_{\hat{r}j}(t)$ for $0 \leqslant t \leqslant 15$ and each $j \in M$, where the additional kernel row of $\hat{r}$ is computed using Eqs. (11) and (12). As expected, these correspond to the steady-state probabilities.

5 Conclusions

We presented a solution to compute a closed-form expression of the equilibrium distribution of MRGPs. The solution leverages the calculus of stochastic state classes, and it can be applied to a given STPN through the implementation in the ORIS tool. In future work, we plan to apply this solution to compute survivability measures [15] for MRGPs; in particular, equilibrium PDFs can be used to characterize the system after a catastrophic failure at the steady state, providing the initial conditions for the transient analysis of system recovery.

A Stochastic Time Petri Nets

STPNs are a formal model of concurrent timed systems where: *transitions* (depicted as vertical bars) represent activities; *places* (depicted as circles) represent discrete components of the logical state, with values encoded by a number of *tokens* (depicted as dots); *directed arcs* from *input* places to transitions and from transitions to *output* places represent token moves triggered by the *firing of transitions*. A transition is enabled when all its input places contain at least one token; its firing removes a token from each input place and adds a token to each output place. The time from the enabling to the firing of a transition is a random variable, and the choice between transitions with equal time to fire is solved by a random switch determined by transition *weights*. Moreover, STPNs can: (1) restrict the enabling of a transition using general constraints on token counts (called *enabling functions*); (2) execute additional updates of token counts after a transition firing (specified by *update functions*); (3) restart selected transitions after a firing (using *reset sets*); (4) impose *priorities* among immediate or deterministic transitions.

Definition 2 (Syntax). *An STPN is a tuple* $\langle P, T, A^-, A^+, B, U, R, EFT, LFT, F, W, Z\rangle$ *where:* P *and* T *are disjoint sets of places and transitions, respectively;* $A^- \subseteq P \times T$ *and* $A^+ \subseteq T \times P$ *are precondition and post-condition relations, respectively;* B, U, *and* R *associate each transition* $t \in T$ *with an enabling function* $B(t)\colon M \to \{\text{TRUE}, \text{FALSE}\}$, *an update function* $U(t)\colon M \to M$, *and a reset set* $R(t) \subseteq T$, *respectively, where* M *is the set of reachable markings* $m\colon P \to \mathbb{N}$*;* EFT *and* LFT *associate each transition* $t \in T$ *with an earliest firing time* $EFT(t) \in \mathbb{Q}_{\geqslant 0}$ *and a latest firing time* $LFT(t) \in \mathbb{Q}_{\geqslant 0} \cup \{\infty\}$ *such that* $EFT(t) \leqslant LFT(t)$*;* F, W, *and* Z *associate each transition* $t \in T$ *with a Cumulative Distribution Function (CDF)* F_t *for its duration* $\tau(t) \in [EFT(t), LFT(t)]$ *(i.e.,* $F_t(x) = P\{\tau(t) \leqslant x\}$, *with* $F_t(x) = 0$ *for* $x < EFT(t)$, $F_t(x) = 1$ *for* $x > LFT(t)$*), a weight* $W(t) \in \mathbb{R}_{>0}$, *and a priority* $Z(t) \in \mathbb{N}$, *respectively.*

A place p is said to be an *input* or *output* place for a transition t if $(p,t) \in A^-$ or $(t,p) \in A^+$, respectively. Following the usual terminology of stochastic Petri nets, a transition t is called *immediate* (IMM) if $EFT(t) = LFT(t) = 0$ and *timed* otherwise; a timed transition is called *exponential* (EXP) if $F_t(x) = 1 - \exp(-\lambda x)$ for some rate $\lambda \in \mathbb{R}_{>0}$, or *general* (GEN) if its time to fire has a non-exponential distribution; as a special case, a GEN transition t is *deterministic* (DET) if $EFT(t) = LFT(t) > 0$. For each transition t with $EFT(t) < LFT(t)$, we assume that F_t can be expressed as the integral function of a probability density function (PDF) f_t, i.e., $F_t(x) = \int_0^x f_t(y)\,dy$. The same notation is also adopted for an IMM or DET transition $t \in T$, which is associated with a Dirac impulse function $f_t(y) = \delta(y - \overline{y})$ with $\overline{y} = EFT(t) = LFT(t)$.

A marking $m \in M$ assigns a natural number of tokens to each place of an STPN. A transition t is *enabled* by m if m assigns at least one token to each of its input places and the enabling function $B(t)(m)$ evaluates to TRUE. The set of transitions enabled by m is denoted as $E(m)$.

Definition 3 (State). *The state of an STPN is a pair* $\langle m, \vec{\tau}\rangle$ *where* $m \in M$ *is a marking and vector* $\vec{\tau} \in \mathbb{R}_{\geqslant 0}^{|E(m)|}$ *assigns a* time to fire $\vec{\tau}(t) \in \mathbb{R}_{\geqslant 0}$ *to each enabled transition* $t \in E(m)$.

Definition 4 (Semantics). *Given an initial marking* m_0*, an* execution *of the STPN is a (finite or infinite) path* $\omega = s_0 \xrightarrow{\gamma_1} s_1 \xrightarrow{\gamma_2} s_2 \xrightarrow{\gamma_3} \cdots$ *such that:* $s_0 = \langle m_0, \vec{\tau}_0\rangle$ *is the initial state, where the time to fire* $\vec{\tau}_0(t)$ *of each enabled transition* $t \in E(m_0)$ *is sampled according to the distribution* F_t*;* $\gamma_i \in T$ *is the* i*th fired transition;* $s_i = \langle m_i, \vec{\tau}_i\rangle$ *is the state reached after the firing of* γ_i*.*
In each state s_i*:*

- *The next transition* γ_{i+1} *is selected from the set of enabled transitions with minimum time to fire and maximum priority according to the distribution given by weights: if* $E_{min} = \arg\min_{t \in E(m_i)} \vec{\tau}_i(t)$ *and* $E_{prio} = \arg\max_{t \in E_{min}} Z(t)$*, then* $t \in E_{prio}$ *is selected with probability* $p_t = W(t) / \left(\sum_{u \in E_{prio}} W(u)\right)$*.*
- *After the firing of* γ_{i+1}*, the new marking* m_{i+1} *is derived by (1) removing a token from each input place of* γ_{i+1}*, (2) adding a token to each output place of* γ_{i+1}*, and (3) applying the update function* $U(\gamma_{i+1})$ *to the resulting marking. A transition* t *enabled by* m_{i+1} *is termed* persistent *if it is distinct from* γ_{i+1}*, it is not contained in* $R(\gamma_{i+1})$*, and it is enabled also by* m_i *and by the intermediate markings after steps (1) and (2); otherwise,* t *is termed* newly enabled *(thus, transitions in the reset set of* γ_{i+1} *are newly enabled if enabled after the firing).*
- *For each newly enabled transition* t*, the time to fire* $\vec{\tau}_{i+1}(t)$ *is sampled according to the distribution* F_t*; for each persistent transition* t*, the time to fire in* s_{i+1} *is reduced by the sojourn time in the previous marking, i.e.,* $\vec{\tau}_{i+1}(t) = \vec{\tau}_i(t) - \vec{\tau}_i(\gamma_{i+1})$*.*

When features are omitted for a transition $t \in T$, default values are assumed as follows: an always-true enabling function $B(t)(m) = $ TRUE; an identity update function $U(t)(m) = m$ for all $m \in M$; an empty reset set $R(t) = \varnothing$; a weight $W(t) = 1$; and, a priority $Z(t) = 0$.

Arc cardinalities greater than 1 can also be introduced in STPN syntax and semantics, letting the firing of a transition remove an arbitrary number of tokens from each input place or add an arbitrary number of tokens to each output place.

References

1. Balsamo, S., Marin, A., Stojic, I.: Perfect Sampling in Stochastic Petri Nets Using Decision Diagrams. In: MASCOTS. pp. 126–135. IEEE Computer Society (2015)
2. Balsamo, S., Marin, A., Stojic, I.: SPNPS: A Tool for Perfect Sampling in Stochastic Petri Nets. In: QEST 2016. LNCS, vol. 9826, pp. 163–166. Springer (2016)
3. Casella, G., Lavine, M., Robert, C.P.: Explaining the perfect sampler. The American Statistician 55(4), 299–305 (2001)
4. Ciardo, G., German, R., Lindemann, C.: A characterization of the stochastic process underlying a stochastic Petri net. IEEE Trans. Softw. Eng. 20(7), 506–515 (Jul 1994)
5. Clark, A., Gilmore, S., Hillston, J., Tribastone, M.: Stochastic Process Algebras. In: SFM 2007. LNCS, vol. 4487, pp. 132–179. Springer (2007)
6. Cloth, L., Haverkort, B.R.: Model checking for survivability. In: QEST 2005. pp. 145–154. IEEE Computer Society (2005)
7. German, R.: Iterative analysis of Markov regenerative models. Perform. Eval. 44(1), 51–72 (2001)
8. Ghasemieh, H., Remke, A., Haverkort, B.R.: Survivability analysis of a sewage treatment facility using hybrid Petri nets. Perform. Eval. 97, 36–56 (2016)
9. Haas, P.: Stochastic Petri Nets. Springer (2002)
10. Horváth, A., Paolieri, M., Ridi, L., Vicario, E.: Transient analysis of non-Markovian models using stochastic state classes. Perform. Eval. 69(7-8), 315–335 (Jul 2012)
11. Kleinrock, L.: Queueing Systems. Wiley (1975)
12. Kulkarni, V.: Modeling and analysis of stochastic systems. Chapman & Hall (1995)
13. Martina, S., Paolieri, M., Papini, T., Vicario, E.: Performance evaluation of Fischer's protocol through steady-state analysis of Markov regenerative processes. In: MASCOTS. pp. 355–360. IEEE Computer Society (2016)
14. Paolieri, M., Biagi, M., Carnevali, L., Vicario, E.: The ORIS Tool: Quantitative Evaluation of Non-Markovian Systems. IEEE Trans. Softw. Eng. 47(6), 1211–1225 (2021)
15. Trivedi, K.S., Xia, R.: Quantification of system survivability. Telecomm. Systems 60(4) (2015)
16. Vicario, E., Sassoli, L., Carnevali, L.: Using stochastic state classes in quantitative evaluation of dense-time reactive systems. IEEE Trans. Softw. Eng. 35(5), 703–719 (Sep/Oct 2009)