

A symbolic approach to quantitative analysis of preemptive real-time systems with non-Markovian temporal parameters

Laura Carnevali, Johnny Giuntini, Enrico Vicario

Dipartimento di Sistemi e Informatica - Università di Firenze

laura.carnevali@unifi.it, giuntini.johnny@gmail.com, enrico.vicario@unifi.it

ABSTRACT

The method of stochastic state classes provides a means for quantitative analysis of a rather wide class of non-Markovian models. As a major and structural limitation, the approach cannot be applied to models encompassing a preemptive policy, which in the practice rules out the mechanism of suspension and resume usually applied in many real-time systems.

We overcome here the limitation by proposing an approach that faces the complexity issues introduced by the suspension/resume mechanism in the structure of supports and distributions of remaining times. In particular, these are distributed over a polyhedral support according to a multivariate joint density function with analytic piecewise form over a partition into polyhedral subdomains. The approach resorts to an imprecise analysis that extends distributions over the tightest DBM zones enclosing polyhedral domains, and approximates them with Bernstein Polynomials to obtain a global (non-piecewise) analytic representation. Computational experience is reported to show the different impact of errors due to the approximation of supports and distributions.

Keywords

Quantitative evaluation, non-Markovian Stochastic Petri Nets, Generalized Semi-Markov Processes, preemptive real-time systems, approximate state space representation, Difference Bounds Matrix, Bernstein Polynomials.

1. INTRODUCTION

Quantitative evaluation of densely-timed models has largely addressed the case of concurrent systems with stochastic temporal parameters, developing on the assumption of behaviors that do not accept suspension and resume. However, this is an essential expressive feature when dealing with real-time systems, which, almost always, run under priority-driven preemptive scheduling [5].

In correctness verification, the case of systems with suspension has been addressed by a few models, which notably include StopWatch

Automata [12], Petri Nets with hyper-arcs [22], Scheduling-TPNs [21], and preemptive Time Petri Nets (pTPNs) [14]. The salient and common trait of this class of models is that the mechanism of suspension changes the order of complexity of the analysis, impairing decidability and polynomial solution time of various problems. To overcome the issue, various approaches resort to an approximation of timing domains. In particular, in [14], an approximation of the state state-space maintains the efficient encoding of Difference Bounds Matrix (DBM) zones and supports exact identification of feasible timings of selected behaviors through a clean-up algorithm, enabling efficient verification of reachability properties under real-time timing constraints.

In quantitative evaluation, the issue of suspension has been addressed by models encompassing the so-called *Preemptive Resume* (PRs) policy [3], also known as *age* policy [17], reaching a limited level of development with respect to models encompassing the more conventional *Preemptive Repeat Different* (PRD) policy [3], also called *enabling memory* policy [17]. In particular, analytical approaches based on a continuous abstraction of time were proposed in [6] for models with exponential and deterministic timers (Deterministic and Stochastic Petri Nets - DSPNs), using Markov Renewal Theory under the so-called enabling restriction that rules out concurrent enabling of multiple generally distributed (GEN) transitions. In [3], the approach is extended to manage a combined use of different preemption policies. A wider extension that relaxes both the enabling restriction limit and accepts any kind of GEN distribution was implemented in the WebSPN Tool [4] through a discrete approximation of time. In [16], a discrete-time variant of Time Petri Nets (TPNs) [10] leverages a maximal step semantics of concurrency to support the representation of preemptive behavior and associate quantitative probabilities with timers and switches. With the development of analysis methods that go beyond the limits of the Markovian assumption and the enabling restriction, quantitative evaluation is increasingly applied to real-time systems. This makes the need to encompass suspension crucial and thus largely emphasizes the relevance of solution techniques that follow the PRs semantics.

In this paper, we propose a symbolic approach to quantitative evaluation of densely-timed preemptive systems with non-Markovian temporal parameters. To this end, we extend the model of stochastic Time Petri Nets (sTPNs) (Section 2) and the method of stochastic state classes [11], [20] (Section 3) to represent resource assignments and to encompass the representation of suspension in the advancement of clocks. In particular, we characterize the complexities of an exact approach, which turns out to be practically impaired

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VALUETOOLS 2011, May 16-20, Paris, France

Copyright © 2011 ICST 978-1-936968-09-1

DOI 10.4108/icst.valuetools.2011.245728

by the fact that timing domains take the form of linear convex polyhedra and state density functions have a piecewise representation over a partition in polyhedral subdomains. We thus propose an imprecise analysis technique that relies on the approximation of both domains and state density functions, obtaining a relevant gain in computational complexity without having a significant impact on performance measures. Computational experience validates the approach on a model of notable complexity by comparing quantitative measures against simulation results (Section 4). As a relevant trait, the approach supports the definition of quantitative metrics to estimate the impact of approximation of both domains and state density functions, which comprises an important step ahead with respect to non-deterministic analysis of preemptive models [14]. Conclusions are finally drawn in Section 5. For the sake of readability, proof of Theorem 3.1 is reported in the Appendix.

2. SPTPN

2.1 Syntax

A *stochastic preemptive Time Petri Net* (spTPN) is a tuple $\langle P; T; A^-; A^+; A'; m_0; EFT^s; LFT^s; \mathcal{F}; \mathcal{C}; Res; Req; Prio \rangle$.

The first 10 elements comprise the model of sTPNs, which are the variant of non-Markovian Stochastic Petri Nets addressed in [11], [20]. P is a set of places. T is a set of transitions disjoint from P . $A^- \subseteq P \times T$, $A^+ \subseteq T \times P$, and $A' \subseteq P \times T$ are sets of precondition, postcondition, and inhibitor arcs (a place p is said to be an *input*, an *output*, or an *inhibitor* place for transition t if $\langle p, t \rangle \in A^-$, $\langle t, p \rangle \in A^+$, or $\langle p, t \rangle \in A'$, respectively). $m_0 : P \rightarrow \mathbb{N}$ is the initial marking associating each place with a non-negative number of tokens. $EFT^s : T \rightarrow \mathbb{R}_0^+$ and $LFT^s : T \rightarrow \mathbb{R}_0^+ \times (\mathbb{R}_0^+ \cup \{\infty\})$ associate each transition with a static *Earliest Firing Time* and a (possibly infinite) static *Latest Firing Time* ($EFT^s(t) \leq LFT^s(t) \forall t \in T$). \mathcal{F} and \mathcal{C} define a measure of probability for non-deterministic choices: $\mathcal{C} : T \rightarrow \mathbb{R}^+$ associates each transition with a positive weight and $\mathcal{F} : T \rightarrow F_t^s(\cdot)$ associates each transition with a static probability distribution supported in the static firing interval $[EFT^s(t), LFT^s(t)]$. If $EFT^s(t) \neq LFT^s(t)$, we assume that $F_t^s(\cdot)$ is absolutely continuous and, thus, that there exists a density-function $f_t^s(\cdot)$ such that $F_t^s(x) = \int_0^x f_t^s(y) dy$.

The last 3 elements extend the model of sTPNs with a mechanism of resource assignment that makes the progress of timed transitions dependent on the availability of a set of preemptable resources. Res is a set of preemptable resources disjoint from P and T . $Req : T \rightarrow 2^{Res}$ and $Prio : T \rightarrow \mathbb{N}$ associate each transition with a subset of Res representing its resource request and with a static priority level, respectively (low priority numbers run first).

2.2 Semantics

The *state* of an spTPN is a pair $\langle m, \tau \rangle$, where $m : P \rightarrow \mathcal{N}$ is a marking and $\tau : T \rightarrow \mathbb{R}_0^+$ associates each transition with a (dynamic) time-to-fire.

Firability. A transition is *enabled* if each of its input places contains at least one token and none of its inhibitor places contains any token. An enabled transition is *progressing* if every resource it requires is not requested by any other enabled transition with a higher level of priority; otherwise, it is *suspended*. A progressing transition t_0 is *firable* if its time-to-fire $\tau(t_0)$ is not higher than that of any other enabled transition. When multiple transitions are firable, the choice is resolved by a random switch determined by \mathcal{C}

$Prob\{t_0 \text{ is selected}\} = \mathcal{C}(t_0) / \sum_{t_i \in T^f(s)} \mathcal{C}(t_i)$, where $T^f(s)$ is the set of transitions that are firable in state s . Since the treatment does not include immediate (IMM) and deterministic (DET) transitions, \mathcal{C} turns out to be irrelevant for the purposes of the analysis [11], but it is essential in the semantics formulation to associate each choice with a measure of probability [9].

Firing. When a transition t_0 fires, the state $\langle m, \tau \rangle$ is replaced by a new state $s' = \langle m', \tau' \rangle$, which we write as $s \xrightarrow{t_0} s'$. Marking m' is derived from m by removing a token from each input place of t_0 and by adding a token to each output place of t_0 :

$$\begin{aligned} m_{t_0 p}(p) &= m(p) - 1 \quad \forall p \cdot \langle p, t_0 \rangle \in A^-, \\ m'(p) &= m_{t_0 p}(p) + 1 \quad \forall p \cdot \langle t_0, p \rangle \in A^+. \end{aligned} \quad (1)$$

Transitions that are enabled both by the intermediate marking $m_{t_0 p}$ and by m' are said *persistent*, while those that are enabled by m' but not by $m_{t_0 p}$ are said *newly-enabled*. If t_0 is still enabled after its own firing, it is always regarded as newly enabled [7], [10].

For any transition t_a newly-enabled after the firing of t_0 , the time-to-fire takes a random value sampled in the static firing interval according to the static probability distribution $F_{t_a}^s(\cdot)$:

$$\begin{aligned} EFT^s(t_a) &\leq \tau'(t_a) \leq LFT^s(t_a), \\ Prob\{\tau'(t_a) \leq x\} &= F_{t_a}^s(x). \end{aligned} \quad (2)$$

For any transition t_i that was progressing in the previous state and is persistent after the firing of t_0 , the time-to-fire is reduced by the time elapsed in the previous state (which is equal to the time-to-fire of t_0 measured at the entrance in the previous state):

$$\tau'(t_i) = \tau(t_i) - \tau(t_0). \quad (3)$$

For any transition t_x that was suspended in the previous state and is persistent after the firing of t_0 , the time-to-fire remains unchanged:

$$\tau'(t_x) = \tau(t_x). \quad (4)$$

3. A SYMBOLIC APPROACH TO THE ANALYSIS OF SPTPN MODELS

The introduction of suspension in the model semantics requires that the analysis method of stochastic state classes be extended. This impacts in relevant manner both on the structure of supports and density functions. We formulate here an exact calculus and propose an approximated approach that makes the analysis viable.

3.1 Stochastic state classes

A *stochastic state class* (stochastic class for short) is a triple $\langle m, D_{\underline{\tau}}, f_{\underline{\tau}} \rangle$ where m is a marking, $\underline{\tau}$ is the vector of times-to-fire of transitions enabled by m , and $f_{\underline{\tau}}$ is a probability density function for $\underline{\tau}$ supported over $D_{\underline{\tau}}$ [11], [20].

DEFINITION 3.1. A *stochastic class* $\Sigma' = \langle m', D'_{\underline{\tau}'}, f'_{\underline{\tau}'}(\cdot) \rangle$ is a *successor of stochastic class* $\Sigma = \langle m, D_{\underline{\tau}}, f_{\underline{\tau}}(\cdot) \rangle$ through t_0 with probability μ , which we write $\Sigma \xrightarrow{t_0, \mu} \Sigma'$, iff the following property holds: if the marking of the net is m and the vector of times to fire of transitions enabled by m is a random variable $\underline{\tau}$ with support $D_{\underline{\tau}}$ and density $f_{\underline{\tau}}(\cdot)$, then the firing of t_0 occurs with probability $\mu > 0$ and leads to a new marking m' and a new vector of times to fire $\underline{\tau}'$ with support $D'_{\underline{\tau}'}$ and density $f'_{\underline{\tau}'}(\cdot)$.

Given an initial stochastic class, the transitive closure of $\xrightarrow{t, \mu}$ defines a *stochastic class graph* (stochastic graph for short), where vertexes are stochastic classes and edges are labeled with a transition t and a probability μ .

3.2 Initial class

We address models where transitions have expolynomial distribution [1], [2] and, furthermore, those with unbounded support are all distributed over $[0, \infty)$ with negative exponential distribution. We assume that the state density function of the initial class is an expolynomial function with global analytic representation over a DBM domain, expressed as the product of two factors separating general and exponential timers:

$$\begin{aligned}
D_{\underline{\tau}} &= D_{\langle \underline{\tau}_G, \underline{\tau}_E \rangle} = D_{\underline{\tau}_G} \times D_{\underline{\tau}_E}, \\
D_{\underline{\tau}_G} &= \begin{cases} \tau_{\beta_{g_1}} - \tau_{\beta_{g_2}} \leq b_{\beta_{g_1} \beta_{g_2}} & (g_1 g_2) \\ \tau_{\beta_{g_1}} - \tau_* \leq b_{\beta_{g_1} *} & (g_1 *) \\ \tau_* - \tau_{\beta_{g_1}} \leq b_{* \beta_{g_1}} & (* g_1) \\ \forall g_1, g_2 \in [0, G-1] \cdot g_1 \neq g_2, \end{cases} \\
D_{\underline{\tau}_E} &= \begin{cases} \tau_* - \tau_{\gamma_e} \leq 0 & (*e) \\ \forall e \in [0, E-1], \end{cases} \\
f_{\underline{\tau}}(\underline{x}) &= f_{\langle \underline{\tau}_G, \underline{\tau}_E \rangle}(\underline{x}_G, \underline{x}_E) = f_{\underline{\tau}_G}(\underline{x}_G) \cdot f_{\underline{\tau}_E}(\underline{x}_E), \\
f_{\underline{\tau}_G}(\underline{x}_G) &= \sum_{z=0}^Z c_z \prod_{g=0}^{G-1} x_{\beta_g}^{\alpha_z \beta_g} e^{-\lambda_z \beta_g x_{\beta_g}}, \\
f_{\underline{\tau}_E}(\underline{x}_E) &= \prod_{e=0}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}, \tag{5}
\end{aligned}$$

where the vector of times-to-fire of enabled transitions is decomposed into two subvectors $\underline{\tau} = \langle \underline{\tau}_G, \underline{\tau}_E \rangle$ encoding the times-to-fire of generally distributed and exponentially distributed transitions, respectively; G and E denote the sizes of $\underline{\tau}_G$ and $\underline{\tau}_E$, respectively; β_g and γ_e denote the indexes of the g -th generally distributed transition and the e -th exponentially distributed transition, respectively (i.e., $\underline{\tau}_G = \langle \tau_{\beta_0}, \dots, \tau_{\beta_{G-1}} \rangle$ and $\underline{\tau}_E = \langle \tau_{\gamma_0}, \dots, \tau_{\gamma_{E-1}} \rangle$); and, τ_* represents the ground time at which the class was entered. The DBM representation has a normal form where coefficients $b_{\beta_{g_1} \beta_{g_2}}$ coincide with the maximum value that can be attained by the difference $\tau_{\beta_{g_1}} - \tau_{\beta_{g_2}}$. The normal form exists uniquely, can be computed in time $O(G^3)$ or even $O(G^2)$, and is univocally identified by the condition $b_{\beta_{g_1} \beta_{g_2}} \leq b_{\beta_{g_1} \beta_{g_3}} + b_{\beta_{g_3} \beta_{g_2}} \forall g_1, g_2, g_3 \in [0, G-1] \cup \{*\}$ with $g_1 \neq g_2 \neq g_3 \neq g_1$ and $\beta_* = *$ [10].

The product form is easily verified in the initial class, where the times-to-fire of all enabled transitions are independent random variables distributed according to their respective static density functions, as if they were newly-enabled. We will illustrate that the property is maintained in all the successor classes.

3.3 Successor detection

The firing of a transition, say a general transition t_{β_0} , is a *possible outgoing event* from a class $\Sigma = \langle m, D_{\underline{\tau}}, f_{\underline{\tau}}() \rangle$ iff t_{β_0} is progressing in Σ and $D_{\underline{\tau}}$ accepts solutions such that t_{β_0} has a time-to-fire not higher than that of any other enabled transition. This occurs iff there is a non-empty set of solutions for the *restricted*

firing domain $D_{t_{\beta_0}}$:

$$\begin{aligned}
D_{t_{\beta_0}} &= D_{\underline{\tau}} \cap \{(\tau_{\beta_0} \leq \tau_{\beta_g}) \wedge (\tau_{\beta_0} \leq \tau_{\gamma_e}), \forall g \in [1, G-1], \\
&\quad \forall e \in [0, E-1] \cdot t_{\beta_g} \text{ and } t_{\gamma_e} \text{ are progressing in } \Sigma\}. \tag{6}
\end{aligned}$$

As in [11], [20], the probability $\mu_{t_{\beta_0}}$ that t_{β_0} is the outgoing event from class Σ is the joint probability that $\underline{\tau}$ belongs to $D_{t_{\beta_0}}$ and t_{β_0} is selected in the random switch among all progressing transitions that share the same time-to-fire of t_{β_0} . Since $F_{\underline{\tau}}^s()$ is absolutely continuous $\forall t \in T$, the probability that t_{β_0} has the same time-to-fire of any other progressing transition is equal to 0. Hence, $\mu_{t_{\beta_0}}$ is equal to the probability that $\underline{\tau}$ belongs to $D_{t_{\beta_0}}$:

$$\mu_{t_{\beta_0}} = \int_{D_{t_{\beta_0}}} f_{\underline{\tau}}(\underline{x}) d\underline{x}. \tag{7}$$

3.4 Successor derivation

The steps of derivation of a successor class require a different development depending on whether the firing transition is generally distributed or exponentially distributed.

3.4.1 Firing of a general transition

Conditioning. The assumption that a general transition t_{β_0} is the next transition to fire conditions $\underline{\tau}$ and yields a new random variable $\underline{\tau}^a = \langle \underline{\tau} \mid (\tau_{\beta_0} \leq \tau_{\beta_i}) \wedge (\tau_{\beta_0} \leq \tau_{\gamma_h}), \forall i \in [1, I-1], \forall h \in [0, H-1] \rangle$, where $\beta_0, \dots, \beta_{I-1}$ and $\gamma_0, \dots, \gamma_{H-1}$ denote the indexes of generally distributed transitions and exponentially distributed transitions, respectively, that are progressing in Σ and persistent in Σ' . The vector $\underline{\tau}^a$ is distributed over:

$$D_{\underline{\tau}^a} = D_{\langle \underline{\tau}_G, \underline{\tau}_E \rangle} \cap \{(\tau_{\beta_0} \leq \tau_{\beta_i}) \wedge (\tau_{\beta_0} \leq \tau_{\gamma_h}), \forall i \in [1, I-1], \forall h \in [0, H-1]\} \tag{8}$$

according to:

$$f_{\underline{\tau}^a}(\underline{x}_G, \underline{x}_E) = \frac{f_{\langle \underline{\tau}_G, \underline{\tau}_E \rangle}(\underline{x}_G, \underline{x}_E)}{\mu_{t_{\beta_0}}}. \tag{9}$$

In particular, if coefficients B denote the maximum value attained by the difference between any two generally distributed timers (including the fictitious ground reference), the normal form of $D_{\underline{\tau}^a}$ can be represented as:

$$D_{\underline{\tau}^a} = \begin{cases} \tau_{\beta_{g_1}}^a - \tau_{\beta_{g_2}}^a \leq B_{\beta_{g_1} \beta_{g_2}} & (g_1 g_2) \\ \tau_{\beta_{g_1}}^a - \tau_*^a \leq B_{\beta_{g_1} *} & (g_1 *) \\ \tau_*^a - \tau_{\beta_{g_1}}^a \leq B_{* \beta_{g_1}} & (* g_1) \\ \tau_{\beta_{g_1}}^a - \tau_{\beta_0}^a \leq B_{\beta_{g_1} \beta_0} & (g_1 0) \\ \tau_{\beta_0}^a - \tau_{\beta_{g_1}}^a \leq B_{\beta_0 \beta_{g_1}} & (0 g_1) \\ \tau_*^a - \tau_{\gamma_h}^a \leq B_{* \beta_0} & (* h) \\ \tau_*^a - \tau_{\gamma_k}^a \leq 0 & (* k) \\ \tau_{\beta_0}^a - \tau_{\gamma_h}^a \leq 0 & (0 h) \\ \forall g_1, g_2 \in [1, G-1], g_1 \neq g_2, \\ \forall h \in [0, H-1], \\ \forall k \in [H, E-1], \end{cases} \tag{10}$$

where: the maximum value attained by $\tau_*^a - \tau_{\gamma_h}^a$ is equal to $B_{* \beta_0}$ since $\tau_*^a - \tau_{\gamma_h}^a = \tau_*^a - \tau_{\beta_0}^a + \tau_{\beta_0}^a - \tau_{\gamma_h}^a \leq \min\{0, B_{* \beta_0}\}$ and $\min\{0, B_{* \beta_0}\} = B_{* \beta_0}$, being $B_{* \beta_0} \leq 0$; the maximum values attained by $\tau_*^a - \tau_{\gamma_k}^a$ and $\tau_{\gamma_k}^a - \tau_{\beta_{g_1}}^a$ are equal to 0 and ∞ , respectively, since constraints added by $D_{\underline{\tau}^a}$ to $D_{\underline{\tau}}$ do not perturb coefficients related to suspended exponentially distributed transitions;

and, constraints on $\tau_{\beta_{g_1}} - \tau_{\gamma_h}$ are not made explicit since they are directly implied by constraints $(g_1 0)$ and $(0h)$, i.e., $\tau_{\beta_{g_1}} - \tau_{\gamma_h} = \tau_{\beta_{g_1}} - \tau_{\beta_0} + \tau_{\beta_0} - \tau_{\gamma_h} \leq B_{\beta_{g_1}\beta_0}$.

Time advancement. At the firing of t_{β_0} , times-to-fire of progressing transitions are reduced by the value of $\tau_{\beta_0}^a$ while those of suspended transitions remain unchanged. This yields a new random variable $\underline{\tau}^b = \langle \tau_{\beta_0}^a, \tau_{\beta_1}^a - \tau_{\beta_0}^a, \dots, \tau_{\beta_{I-1}}^a - \tau_{\beta_0}^a, \tau_{\beta_I}^a, \dots, \tau_{\beta_{G-1}}^a - \tau_{\beta_0}^a, \tau_{\gamma_0}^a, \dots, \tau_{\gamma_{H-1}}^a - \tau_{\beta_0}^a, \tau_{\gamma_H}^a, \dots, \tau_{\gamma_{E-1}}^a \rangle$ distributed over a domain $D_{\underline{\tau}^b}$ that can be represented as the product of two factors separating generally distributed and exponential timers:

$$D_{\underline{\tau}^b} = D_{\underline{\tau}_G^b} \times D_{\underline{\tau}_E^b}, \quad (11)$$

where $D_{\underline{\tau}_G^b}$ is made by DBM constraints and $D_{\underline{\tau}_E^b} = [0, \infty)^{E-1}$.

The density function $f_{\underline{\tau}^b}()$ of $\underline{\tau}^b$ is obtained from $f_{\underline{\tau}^a}()$ by shifting the components of progressing transitions by x_{β_0} :

$$\begin{aligned} f_{\underline{\tau}^b}(\underline{x}_G, \underline{x}_E) &= f_{\underline{\tau}^a}(x_{\beta_0}, x_{\beta_1} + x_{\beta_0}, \dots, x_{\beta_{I-1}} + x_{\beta_0}, x_{\beta_I}, \\ &\quad \dots, x_{\beta_{G-1}}, x_{\gamma_0} + x_{\beta_0}, \dots, x_{\gamma_{H-1}} + x_{\beta_0}, x_{\gamma_H}, \dots, x_{\gamma_E}) \\ &= \frac{f_{\underline{\tau}_G^a}(x_{\beta_0}, x_{\beta_1} + x_{\beta_0}, \dots, x_{\beta_{I-1}} + x_{\beta_0}, x_{\beta_I}, \dots, x_{\beta_{G-1}})}{\prod_{h=0}^{H-1} e^{-\lambda_{\gamma_h} x_{\beta_0}} \cdot \prod_{e=0}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}}. \end{aligned} \quad (12)$$

According to this, the density function of $\underline{\tau}^b$ can be expressed as the product of two factors $f_{\underline{\tau}^b}(\underline{x}_G, \underline{x}_E) = f_{\underline{\tau}_G^b}(\underline{x}_G) \cdot f_{\underline{\tau}_E^b}(\underline{x}_E)$ separating generally distributed and exponential timers:

$$\begin{aligned} f_{\underline{\tau}_G^b}(\underline{x}_G) &= \prod_{h=0}^{H-1} e^{-\lambda_{\gamma_h} x_{\beta_0}} \\ &\quad \cdot \frac{f_{\underline{\tau}_G^a}(x_{\beta_0}, x_{\beta_1} + x_{\beta_0}, \dots, x_{\beta_{I-1}} + x_{\beta_0}, x_{\beta_I}, \dots, x_{\beta_{G-1}})}{\prod_{h=0}^{H-1} e^{-\lambda_{\gamma_h} x_{\beta_0}}}, \\ f_{\underline{\tau}_E^b}(\underline{x}_E) &= \prod_{e=0}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}. \end{aligned} \quad (13)$$

The time-to-fire of t_{β_0} is then eliminated from $\underline{\tau}^b$, yielding a new random variable $\underline{\tau}^c = \langle \tau_{\beta_1}^b, \dots, \tau_{\beta_{G-1}}^b, \tau_{\gamma_0}^b, \dots, \tau_{\gamma_{E-1}}^b \rangle$, which is distributed over the projection of $D_{\underline{\tau}^b}$ that eliminates τ_{β_0} :

$$D_{\underline{\tau}^c} = D_{\underline{\tau}^b} \downarrow_{\beta_0} = (D_{\underline{\tau}_G^b} \downarrow_{\beta_0}) \times D_{\underline{\tau}_E^b} \quad (14)$$

according to:

$$\begin{aligned} f_{\underline{\tau}^c}(\underline{x}_G^c, \underline{x}_E) &= \int_{D_{\underline{\tau}^b}} f_{\underline{\tau}^b}(\underline{x}_G, \underline{x}_E) dx_{\beta_0} \\ &= \int_{D_{\underline{\tau}_G^b}^{\beta_0}(\underline{x}_G^c)} f_{\underline{\tau}_G^b}(\underline{x}_G) dx_{\beta_0} \cdot f_{\underline{\tau}_E^b}(\underline{x}_E), \end{aligned} \quad (15)$$

where $\underline{x}_G^c = \langle x_{\beta_1}, \dots, x_{\beta_{G-1}} \rangle$; $D_{\underline{\tau}^b} \downarrow_{\beta_0} \stackrel{def}{=} \{ \underline{x}_G^c \in \mathbb{R}^{G-1} \mid \exists x_{\beta_0} \in \mathbb{R} \text{ such that } \langle x_{\beta_0}, \underline{x}_G^c \rangle \in D_{\underline{\tau}^b} \}$; and, the support $D_{\underline{\tau}_G^b}^{\beta_0}(\underline{x}_G^c) \stackrel{def}{=} \{ x_{\beta_0} \in \mathbb{R} \mid \exists \underline{x}_G^c \in \mathbb{R}^{G-1} \text{ such that } \langle x_{\beta_0}, \underline{x}_G^c \rangle \in D_{\underline{\tau}^b} \}$ represents all possible values of τ_{β_0} . According to this, the subvector $\underline{\tau}_G^c = \langle \tau_{\beta_1}^c, \dots, \tau_{\beta_{G-1}}^c \rangle$ is distributed over $D_{\underline{\tau}_G^c} = D_{\underline{\tau}_G^b} \downarrow_{\beta_0}$ according to $f_{\underline{\tau}_G^c}(\underline{x}_G^c) = \int_{D_{\underline{\tau}_G^b}^{\beta_0}(\underline{x}_G^c)} f_{\underline{\tau}_G^b}(\underline{x}_G) dx_{\beta_0}$ and the subvector $\underline{\tau}_E^c = \langle \tau_{\gamma_0}^c, \dots, \tau_{\gamma_{E-1}}^c \rangle$ is distributed over $D_{\underline{\tau}_E^c} = [0, \infty)^{E-1}$

according to $f_{\underline{\tau}_E^c}(\underline{x}_E) = f_{\underline{\tau}_E^b}(\underline{x}_E)$, with $D_{\underline{\tau}^c} = D_{\underline{\tau}_G^c} \times D_{\underline{\tau}_E^c}$ and $f_{\underline{\tau}^c}(\underline{x}) = f_{\underline{\tau}_G^c}(\underline{x}_G) \cdot f_{\underline{\tau}_E^c}(\underline{x}_E)$.

Disabling. When t_{β_0} fires, one or more transitions may be disabled. If a generally distributed transition, say $t_{\beta_{G-1}}$, is disabled, its elimination yields a new vector of times-to-fire $\underline{\tau}^d = \langle \tau_{\beta_1}^c, \dots, \tau_{\beta_{G-2}}^c, \underline{\tau}_E^c \rangle$, distributed over:

$$D_{\underline{\tau}^d} = (D_{\underline{\tau}_G^c} \downarrow_{\beta_{G-1}}) \times D_{\underline{\tau}_E^c} \quad (16)$$

according to:

$$f_{\underline{\tau}^d}(\underline{x}_G^d, \underline{x}_E^c) = \int_{D_{\underline{\tau}_G^c}} f_{\underline{\tau}_G^c}(\underline{x}_G^c) dx_{\beta_{G-1}} \cdot f_{\underline{\tau}_E^c}(\underline{x}_E^c), \quad (17)$$

where $\underline{x}_G^d = \langle x_{\beta_1}, \dots, x_{\beta_{G-2}} \rangle$. In a similar manner, the elimination of an exponentially distributed transition, say $t_{\gamma_{E-1}}$, yields the new random variable $\underline{\tau}^d = \langle \tau_{\gamma_0}^c, \dots, \tau_{\gamma_{E-2}}^c \rangle$, distributed over $D_{\underline{\tau}^d} = D_{\underline{\tau}_G^c} \times [0, \infty)^{E-2}$ according to $f_{\underline{\tau}^d}(\underline{x}_G^c, \underline{x}_E^d) = f_{\underline{\tau}_G^c}(\underline{x}_G^c) \cdot \prod_{e=0}^{E-2} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}$, where $\underline{x}_E^d = \langle x_{\gamma_0}, \dots, x_{\gamma_{E-2}} \rangle$. When multiple generally distributed or exponentially distributed transitions are disabled, the step is repeated for each of them.

As in [20], the step of disabling may partition the subdomain of generally distributed timers into a finite set of DBM subdomains, over each of which the state density function accepts an analytic representation. To resort to the representation of Eq.(5), the state density function is approximated through Bernstein Polynomials [18], [23] as it will be illustrated in Section 3.5.

Newly enabling. If transitions $t_{\delta_0}, \dots, t_{\delta_{N-1}}$ are newly enabled, the vector $\underline{\tau}' = \langle \tau_{\delta_0}, \dots, \tau_{\delta_{N-1}} \rangle$ of times-to-fire of transitions enabled in the destination class Σ' is distributed over:

$$D' = D_{\underline{\tau}^d} \times \prod_{i=0}^{N-1} [EFT^s(t_{\delta_i}), LFT^s(t_{\delta_i})] \quad (18)$$

according to:

$$f_{\underline{\tau}'}(\underline{x}_G^d, \underline{x}_E^d, x_{\delta_0}, \dots, x_{\delta_{N-1}}) = f_{\underline{\tau}^d}(\underline{x}_G^d, \underline{x}_E^d) \cdot \prod_{i=0}^{N-1} f_{t_{\delta_i}^s}(x_{\delta_i}), \quad (19)$$

preserving the product decomposition of state density functions. Note that this is the step where the form of state density functions is determined by static density functions associated with transitions in the model.

3.4.2 Firing of an exponential transition

The assumption that the firing transition is an exponential transition, say t_{γ_0} , gives rise to two subcases.

If no general transition is progressing, then the steps of conditioning and time advancement eliminate the factor pertaining to t_{γ_0} from both the class domain and the state density function, without involving the other persistent transitions. In fact, the assumption that t_{γ_0} is the next transition to fire conditions the vector of times-to-fire $\underline{\tau}$ and yields a new random variable $\underline{\tau}^a = \langle \tau_{\gamma_0}^a \leq \tau_{\gamma_h}, \forall h \in [1, H-1] \rangle$, distributed over $D_{\underline{\tau}^a} = (D_{\underline{\tau}_G} \times [0, \infty)^E) \cap \{ \tau_{\gamma_0} \leq \tau_{\gamma_h}, \forall h \in [1, H-1] \}$ according to $f_{\underline{\tau}^a}(\underline{x}_G, \underline{x}_E) = f_{\underline{\tau}_G}(\underline{x}_G) \cdot f_{\underline{\tau}_E}(\underline{x}_E) / \mu_{\gamma_0} = f_{\underline{\tau}_G}(\underline{x}_G) \cdot \prod_{e=0}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}} \cdot \Lambda / \lambda_{\gamma_0}$, where $\Lambda = \sum_{h=0}^{H-1} \lambda_{\gamma_h}$. The step of time advancement yields a

new vector of times-to-fire $\underline{\tau}^b = \langle \tau_{\beta_I}^a, \dots, \tau_{\beta_{G-1}}^a, \tau_{\gamma_0}^a, \tau_{\gamma_1}^a - \tau_{\gamma_0}^a, \dots, \tau_{\gamma_{H-1}}^a - \tau_{\gamma_0}^a, \tau_{\gamma_H}^a, \dots, \tau_{\gamma_{E-1}}^a \rangle$, distributed over $D_{\underline{\tau}^b} = D_{\underline{\tau}^a}$ according to $f_{\underline{\tau}^b}(\underline{x}_G, \underline{x}_E) = f_{\underline{\tau}^a}(\underline{x}_G, x_{\gamma_0}, x_{\gamma_1} + x_{\gamma_0}, \dots, x_{\gamma_{H-1}} + x_{\gamma_0}, x_{\gamma_H}, \dots, x_{\gamma_{E-1}}) = f_{\underline{\tau}^a}(\underline{x}_G) \cdot \Lambda / \lambda_{\gamma_0} \cdot \lambda_{\gamma_0} e^{-\Lambda x_{\gamma_0}} \cdot \prod_{e=1}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}} = f_{\underline{\tau}^a}(\underline{x}_G) \cdot \Lambda \cdot e^{-\Lambda x_{\gamma_0}} \cdot \prod_{e=1}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}$. The projection that eliminates t_{γ_0} yields the vector of times-to-fire $\underline{\tau}^c = \langle \tau_{\beta_I}^b, \dots, \tau_{\beta_{G-1}}^b, \tau_{\gamma_1}^b, \dots, \tau_{\gamma_{E-1}}^b \rangle$, distributed over $D_{\underline{\tau}^c} = D_{\underline{\tau}^G} \times [0, \infty)^{E-1}$ according to $f_{\underline{\tau}^c}(\underline{x}_G, x_{\gamma_1}, \dots, x_{\gamma_{E-1}}) = f_{\underline{\tau}^a}(\underline{x}_G) \cdot \prod_{e=1}^{E-1} \lambda_{\gamma_e} e^{-\lambda_{\gamma_e} x_{\gamma_e}}$. Finally, the elimination of disabled transitions and the addition of newly enabled transitions is performed as described in Section 3.4.1.

If at least a general transition is progressing, then the conditioning of the vector of times-to-fire $\underline{\tau}$ yields a new random variable $\underline{\tau}^a = \langle \tau_{\beta_I}, \dots, \tau_{\beta_{G-1}}, \tau_{\gamma_0}, \dots, \tau_{\gamma_{H-1}} \rangle$, distributed over $D_{\underline{\tau}^a} = D_{\langle \tau_{\beta_I}, \dots, \tau_{\beta_{G-1}}, \tau_{\gamma_0}, \dots, \tau_{\gamma_{H-1}} \rangle} \cap \{(\tau_{\gamma_0} \leq \tau_{\beta_i}) \wedge (\tau_{\gamma_0} \leq \tau_{\gamma_h}), \forall i \in [0, I-1], \forall h \in [1, H-1]\}$, distributed over $D_{\underline{\tau}^a} = D_{\langle \tau_{\beta_I}, \dots, \tau_{\beta_{G-1}}, \tau_{\gamma_0}, \dots, \tau_{\gamma_{H-1}} \rangle} \cap \{(\tau_{\gamma_0} \leq \tau_{\beta_i}) \wedge (\tau_{\gamma_0} \leq \tau_{\gamma_h}), \forall i \in [0, I-1], \forall h \in [1, H-1]\}$ according to $f_{\underline{\tau}^a}(\underline{x}_G, \underline{x}_E) = f_{\langle \tau_{\beta_I}, \dots, \tau_{\beta_{G-1}}, \tau_{\gamma_0}, \dots, \tau_{\gamma_{H-1}} \rangle}(\underline{x}_G, \underline{x}_E) / \mu_{t_{\gamma_0}}$. In so doing, τ_{γ_0} turns out to be distributed over $[0, \min_{i \in [0, H-1]} B_{\beta_i}]$ according to $f_{\tau_{\gamma_0}}(x_{\gamma_0}) = \lambda_{\gamma_0} e^{-\lambda_{\gamma_0} x_{\gamma_0}} / \mu_{t_{\gamma_0}}$ and, thus, t_{γ_0} can be regarded as a general transition. According to this, t_{γ_0} is inserted in the set of general transitions and successor derivation proceeds from the step of time advancement described in Section 3.4.1 through the steps of disabling and newly enabling.

3.5 Successor approximation

The step of time advancement described in Section 3.4.1 subtends notable complexities both in the derivation of the firing domain and in the computation of the state density function. We characterize here the complexities of exact analysis and propose an imprecise approach that approximates both domains and state density functions, obtaining a relevant gain in computational complexity without significantly affecting the accuracy of performance measures.

3.5.1 Approximation of class domains

In the derivation of the timing domain, the elimination of τ_{β_0} from $D_{\underline{\tau}^b}$ yields a domain $D_{\underline{\tau}^c}$ [14] which includes constraints in DBM form but also linear constraints with more than two unknown values. According to this, $D_{\underline{\tau}^c}$ is not in DBM form and, thus, the space of DBM firing domains is not closed with respect to the succession relation induced by the semantics of sTPNs. In particular, $D_{\underline{\tau}^c}$ takes the form of a linear convex polyhedron which gets more and more complex as the succession transformation is repeatedly applied, yielding a number of inequalities which is exponential in the number of generally distributed enabled transitions. According to this, the derivation of the exact form of $D_{\underline{\tau}^c}$ becomes a general linear programming problem [13], which can be solved with at least polynomial complexity in the number of domain inequalities and, thus, exponential complexity in the number of generally distributed enabled transitions.

We overcome space and time complexities of representation and manipulation of polyhedral constraints by replacing the subdomain of general transitions with its approximation $\tilde{D}_{\underline{\tau}^c}$ that discards inequalities that are not in DBM form [14]. Since all inequalities of $\tilde{D}_{\underline{\tau}^c}$ are included in $D_{\underline{\tau}^c}$, $\tilde{D}_{\underline{\tau}^c}$ admits any solution admitted by $D_{\underline{\tau}^c}$ plus additional solutions that would not be feasible for the exact representation of the firing domain, which we call *false behaviors*. In [14], $\tilde{D}_{\underline{\tau}^c}$ is proven to be the tightest possible embedding DBM of $D_{\underline{\tau}^c}$ and an algorithm is provided for the clean-up of false behaviors and the tight estimation of the timing profile of feasible

behaviors. In particular, $D_{\underline{\tau}^c}$ is actually a linear convex polyhedron iff the parent class includes both progressing and suspended transitions that are persistent after the firing of a transition taking a non-deterministic time-to-fire; otherwise, it preserves the DBM representation and does not require any approximation. This enables straightforward identification of classes where approximation errors are actually introduced and supports techniques that leverage structural properties of a model to confine the effects of approximation.

3.5.2 Approximation of state density functions

In the derivation of the state density function, the support $D_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)$ of Eq.(15) is an interval $D_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = [E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c), L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)]$, with:

$$E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \max_{\substack{i \in [1, I-1] \\ x \in [I, G-1]}} \{-B_{*\beta_0}, -B_{\beta_x \beta_i} + x_{\beta_x} - x_{\beta_i}, -B_{*\beta_i} - x_{\beta_i}, -B_{\beta_x \beta_0} + x_{\beta_x}\} \quad (a)$$

$$L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \min_{\substack{i \in [1, I-1] \\ x \in [I, G-1]}} \{B_{\beta_0 *}, B_{\beta_i \beta_x} + x_{\beta_x} - x_{\beta_i}, B_{\beta_i *}, -x_{\beta_i}, B_{\beta_0 \beta_x} + x_{\beta_x}\} \quad (20)$$

Eq.(20) partitions $D_{\underline{\tau}^c}$ into a finite set of polyhedral subdomains within each of which $E_{\underline{\tau}^c}^{\beta_0}()$ and $L_{\underline{\tau}^c}^{\beta_0}()$ accept analytic representation. The density function $f_{\underline{\tau}^c}()$ accepts a continuous piecewise representation over this partition, since static density functions of transitions in the model are continuous functions and the steps of conditioning and time-advancement preserve the property of continuity of state density functions. The derivation of the exact representation of $f_{\underline{\tau}^c}()$ thus requires exponential complexity in the number of generally distributed enabled transitions.

Theorem 3.1. *Integration bounds $E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)$ and $L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)$ are piecewise continuous functions with analytic representation over each element of a partition of $D_{\underline{\tau}^c}$ in polyhedral subdomains R_{ixjy} , with $i, j \in [0, I-1]$ and $x, y \in [I, G-1]$. Within each subdomain R_{ixjy} , the two integration bounds are either constants, or mono-variate linear functions of a single component of the vector \underline{x}_G^c , or bivariate linear functions of two components of the vector \underline{x}_G^c . In particular, if $\underline{x}_G^c \in R_{ixjy}$, then:*

$$E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \begin{cases} -B_{*\beta_0} & \text{iff } i = 0, x = I \\ -B_{*\beta_i} - x_{\beta_i} & \text{iff } i > 0, x = I \\ -B_{\beta_x \beta_0} + x_{\beta_x} & \text{iff } i = 0, x > I \\ -B_{\beta_x \beta_i} + x_{\beta_x} - x_{\beta_i} & \text{iff } i > 0, x > I \end{cases}$$

$$L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \begin{cases} B_{\beta_0 *}, & \text{iff } j = 0, y = I \\ B_{\beta_j *}, -x_{\beta_j} & \text{iff } j > 0, y = I \\ B_{\beta_0 \beta_y} + x_{\beta_y} & \text{iff } j = 0, y > I \\ B_{\beta_j \beta_y} + x_{\beta_y} - x_{\beta_j} & \text{iff } j > 0, y > I \end{cases} \quad (21)$$

Note that, when domain $D_{\underline{\tau}^c}$ maintains the DBM representation, the integration bounds $E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)$ and $L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c)$ become:

$$E_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \max_{i \in [1, G-1]} \{-B_{*\beta_0}, -B_{*\beta_i} - x_{\beta_i}\},$$

$$L_{\underline{\tau}^c}^{\beta_0}(\underline{x}_G^c) = \min_{i \in [1, G-1]} \{B_{\beta_0 *}, B_{\beta_i *}, -x_{\beta_i}\}, \quad (22)$$

which partitions $D_{\underline{\tau}^c}$ into a finite set of DBM subdomains R_{ij} with $i, j \in [1, G-1]$, within each of which $E_{\underline{\tau}^c}^{\beta_0}()$ and $L_{\underline{\tau}^c}^{\beta_0}()$ are

either constants or monivariate functions of a single component of the vector \underline{x}_G^c [20]. In particular, if $\underline{x}_G^c \in R_{ij}$, then:

$$\begin{aligned} E_{\underline{x}_G^c}^{\beta_0}(\underline{x}_G) &= \begin{cases} -B_{*0} & \text{iff } i = 0 \\ -B_{0i} + x_i & \text{iff } i > 0 \end{cases} \\ L_{\underline{x}_G^c}^{\beta_0}(\underline{x}_G) &= \begin{cases} B_{0*} & \text{iff } j = 0 \\ B_{j0} + x_j & \text{iff } j > 0 \end{cases} \end{aligned} \quad (23)$$

As a consequence of this, the density function $f_{\underline{x}_G^c}(\cdot)$ accepts piecewise representation over a partition of $D_{\underline{x}_G^c}$ in DBM subdomains.

We cope with the complexity of derivation of state density functions by employing global approximant functions. Since $f_{\underline{x}_G^c}(\cdot)$ is a continuous function over the bounded and thus compact support $D_{\underline{x}_G^c}$, according to Weierstrass theorem $f_{\underline{x}_G^c}(\cdot)$ is bounded and can be arbitrarily well approximated by polynomials.

We employ multivariate Bernstein Polynomials [18], [23], which were successfully used in the approximation of state density functions of models that do not encompass preemptive behavior [20]. Bernstein Polynomials approximate a function defined over a compact hyper-rectangle by weighting a kernel of multivariate monomials according to the samples of the approximated function taken over a regular grid. If b_{ij} denote the coefficients of the normal form of $\tilde{D}_{\underline{x}_G^c}$, then the Bernstein approximant $\tilde{f}_{\underline{x}_G^c}(\cdot)$ of $f_{\underline{x}_G^c}(\cdot)$ that is defined over the minimum embedding hyper-rectangle $\tilde{R}_{\underline{x}_G^c} = \prod_{g=1}^{G-1} [-b_{*\beta_g}, b_{\beta_g*}]$ of $\tilde{D}_{\underline{x}_G^c}$ and takes K_g samples for each variable x_{β_g} is:

$$\begin{aligned} \tilde{f}_{\underline{x}_G^c}(\underline{x}_G^c) &= \sum_{\substack{k_g \in [0, K_g - 1] \\ g \in [1, G - 1]}} f_{\underline{x}_G^c} \left(-b_{*\beta_1} + \frac{k_1(b_{\beta_1*} - (-b_{*\beta_1}))}{K_1 - 1}, \dots, \right. \\ &\quad \left. -b_{*\beta_{G-1}} + \frac{k_{G-1}(b_{\beta_{G-1}*} - (-b_{*\beta_{G-1}}))}{K_{G-1} - 1} \right) \\ &\cdot \prod_{g=1}^{G-1} \binom{K_g - 1}{k_g} \frac{(x_{\beta_g} - (-b_{*\beta_g}))^{k_g} (b_{\beta_g*} - x_{\beta_g})^{K_g - 1 - k_g}}{(b_{\beta_g*} - (-b_{*\beta_g}))^{K_g - 1}}, \end{aligned} \quad (24)$$

where $\sum_{k_g \in [0, K_g - 1], g \in [1, G - 1]}$ denotes the multiple summation $\sum_{k_1=0}^{K_1-1} \sum_{k_2=0}^{K_2-1} \dots \sum_{k_{G-1}=0}^{K_{G-1}-1}$ and $f_{\underline{x}_G^c}(\cdot)$ is extended over the entire hyper-rectangle by assigning null value to samples belonging to $\tilde{R}_{\underline{x}_G^c}$ but not to $D_{\underline{x}_G^c}$, i.e., $f_{\underline{x}_G^c}(\underline{y}) = 0$ if $\underline{y} \in \tilde{R}_{\underline{x}_G^c} \setminus D_{\underline{x}_G^c}$. Note that samples of $f_{\underline{x}_G^c}(\cdot)$ belonging to $D_{\underline{x}_G^c}$ are derived in a straightforward manner by resolving Eq.(15), since, for a given $\underline{y} \in D_{\underline{x}_G^c}$, the integrand function $f_{\underline{x}_G^c}(x_{\beta_0}, \underline{y})$ is a monivariate function of x_{β_0} and the integration bounds $E_{\underline{x}_G^c}^{\beta_0}(\underline{y})$ and $L_{\underline{x}_G^c}^{\beta_0}(\underline{y})$ are constant values (obtained as the maximum and the minimum value of the two sets defined in Eq.(20), respectively, whose elements are all constant for assigned values of variables $x_{\beta_1}, \dots, x_{\beta_{G-1}}$).

On the one hand, approximation based on Bernstein Polynomials exhibits a set of favorable properties that nicely fit the needs of our application context: the approximant is *global*, in the sense that it has analytic representation over the entire domain of the approximated function; the approximant is *positive*, since all polynomials in the kernel and the samples of state density functions are positive; the approximant is *simple to derive*, since it is obtained in a straightforward manner from the samples of the approximated function; and, the approximant *converges* uniformly to the approximated function as the number of samples is increased, provided that the approximated function is continuous (which is the case of our state density functions), with an approximation error bounded by a

Lipschitz inequality [18]. On the other hand, Bernstein Polynomials do not preserve the integral of the approximated function. For this reason, the approximant function is normalized with respect to its own integral over $\tilde{D}_{\underline{x}_G^c}$ to guarantee unit-measure.

3.6 Equivalence between stochastic classes

As in [20], the test of equivalence between state density functions is relaxed by assuming that two classes $\Sigma = \langle m, D_{\underline{x}}, f_{\underline{x}}(\cdot) \rangle$ and $\Sigma' = \langle m', D'_{\underline{x}'}, f'_{\underline{x}'}(\cdot) \rangle$, with $f_{\underline{x}} = f_{\underline{x}_G} \cdot f_{\underline{x}_E}$ and $f'_{\underline{x}'} = f'_{\underline{x}_G} \cdot f'_{\underline{x}_E}$, are equivalent as soon as *i*) they have the same underlying state class (i.e., $m = m'$ and $D_{\underline{x}} = D'_{\underline{x}'}$); *ii*) they have the same rate of exponential transitions (i.e., $\lambda_{\gamma_i} = \lambda'_{\gamma_i} \forall i \in [0, E - 1]$); and *iii*) the distance between $f_{\underline{x}_G}(\cdot)$ and $f'_{\underline{x}'_G}(\cdot)$ measured according to a metrics $\|\cdot\|_d$ is below a given threshold of tolerance δ (i.e., $\|f_{\underline{x}_G}(\cdot) - f'_{\underline{x}'_G}(\cdot)\|_d = \frac{1}{D_{\underline{x}_G}} \int_{D_{\underline{x}_G}} |f_{\underline{x}_G}(\underline{x}_G) - f'_{\underline{x}'_G}(\underline{x}_G)| d\underline{x}_G \leq \delta$). The metrics $\|\cdot\|_d$ is approximated in discrete form by evaluating it in correspondence with samples taken over the regular grid of Bernstein approximants and taking into account border effects [20]:

$$\|f_{\underline{x}}(\cdot)\|_d \simeq \frac{\sum_{\substack{0 \leq k_g \leq K_g - 1 \\ g \in [0, G - 1]}} A(k_0, \dots, k_{G-1}) \cdot \left| f_{\underline{x}}\left(\frac{k_0}{K_0}, \dots, \frac{k_{G-1}}{K_{G-1}}\right) \right|}{\sum_{\substack{0 \leq k_g \leq K_g - 1 \\ g \in [0, G - 1]}} A(k_0, \dots, k_{G-1})} \quad (25)$$

where $A(k_0, \dots, k_{G-1}) = 1/2^{\text{border}(k_0, \dots, k_{G-1})}$ and $\text{border}(k_0, \dots, k_{G-1})$ denotes the number of elements of $\langle k_0, \dots, k_{G-1} \rangle$ that are equal to 0 or K_g .

4. COMPUTATIONAL EXPERIENCE

The accuracy maintained in the enumeration of the stochastic class graph was evaluated by estimating the impact of errors due to the approximation of domains and state density functions, respectively. On the one hand, the error due to the replacement of polyhedral subdomains of general transitions with their tightest embedding DBM was estimated through the portion of points over a regular grid that belong to the DBM zone but not to the polyhedron, which comprises a qualitative measure of false behaviors introduced in the approximation of domains. Detected false behaviors were then weighted according to the samples of the approximated state density function taken over the grid, yielding a quantitative measure of false behaviors introduced in the approximation of domains. We derived the number of classes including false behaviors, together with the average and the maximum values of the percentage and the probability of false behaviors in one of these classes. On the other hand, the error due to the approximation of the factor of state density functions associated with general transitions was evaluated by estimating the distance between approximated functions and their approximant functions through the metrics $\|\cdot\|_d$. We derived the number of classes with non-null d-distance together with the average and the maximum value of the d-distance in one of these classes, both after the step of time-advancement and after the step of disabling. All accuracy metrics were evaluated over a regular grid that takes 10 samples for each variable associated with a general transition.

The overall approach was then validated by evaluating steady state probabilities of reachable markings in the Discrete Time Markov Chain embedded in the stochastic class graph [11], [20] and by comparing them against simulation results. In particular, we com-

pared the overall fit, together with the average and maximum difference in the probability of the same marking. Both analysis and simulation results were obtained through a preliminary implementation in the Oris Tool [15].

In Sects.4.1 and 4.2, we discuss the application of approximated analysis on two examples: the former includes a low number of markings but is still sufficient to illustrate the complexities of the approach, while the latter is a model of notable complexity.

4.1 The M/G/2/3/3 queue

Fig. 1 shows the model of an M/G/2/3/3 queue. The system is a closed queue with 3 customers and 2 servers: transitions t_{10} , t_{20} , and t_{30} model the arrivals of customer 1, 2, and 3, respectively, and have exponentially distributed firing times with rate 0.003; transitions t_{11} , t_{21} , and t_{31} represent the completion of jobs of customer 1, 2, and 3, respectively, by server 1, 1, and 2, respectively, and they have uniformly distributed firing times on the interval $[100, 200]$. The queue has a preemptive service in such a way that a job of customer 2 is preempted as soon as a new job of customer 1 eventually arrives at the server, and it is resumed as soon as the server becomes available again, i.e., the service time of the recovered job is equal to the residual service time of the preempted job. This is obtained by associating t_{10} and t_{20} with the same resource request, i.e., t_{10} and t_{20} require resource r_1 , but with different priority levels, i.e., t_{10} and t_{20} have priority level 1 and 2, respectively.




Figure 1: An sTPN representing an M/G/2/3/3 queue.

In general, the semantics of sTPNs combines the *Preemptive Repeat Different* (PRD) and the *Preemptive Resume* (PRs) policies of [3], [6]: when a transition is disabled by the lack of a token in any input place, its time-to-fire is reset; when a transition is suspended by the lack of any required resource, its time-to-fire is maintained and resumed when the transition is assigned the resource again. In the formulation of [3], [6], instead, the clock of a transition disabled by the lack of a token is reset or maintained if the transition is associated with PRD or PRs policy, respectively. With respect to the model of [3], [6], sTPNs make the progress of times to fire dependent on both the presence of tokens into input places and on the availability of preemptable resources, enabling separate representation of inter-process communication mechanisms from real-time concurrency on resources. This impacts on modeling convenience, facilitating the representation of task-sets [8], [19] as usually encountered in the practice of real-time systems. To illustrate the concept and facilitate comparison, Fig. 2 shows the model of the M/G/2/3/3 queue using the semantics and the graphical notation from [4], [7]. Note that in this case the inhibitor arc from place p_{11} to transition t_{21} accounts for the preemption. Also note that the ex-

ample extends the case of the preemptive M/D/1/2/2 queue of [6] in two aspects: DET timers are replaced with a GEN distribution with non-pointlike support; the number of servers and customers is extended to 2 and 3, respectively, so as to stress the complexity of the analysis by including classes with polyhedral domain.

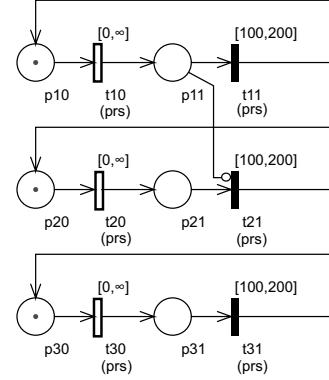


Figure 2: The M/G/2/3/3 queue of Fig. 1 represented according to the formulation of [3].

Approximated analysis with threshold δ equal to 0.0001 and Bernstein degree equal to 4 enumerates 68 stochastic classes, covered by 8 markings, in nearly 10 seconds. Approximation of timing domains may affect any class where t_{10} , t_{20} , and t_{30} are concurrently enabled (since t_{10} and t_{30} are progressing, while t_{20} is suspended) and, in particular, this occurs in 3 classes, as shown in Table 1. Although the portion of false behaviors in any of these classes seems to be relevant, with average and maximum percentage equal to 17.01 % and 25.44 %, respectively, their probability is extremely low, with average and maximum values equal to 0.0124 and 0.0180, respectively.

Metrics on domains	average	max	# classes
% false behaviors	17.01 %	25.44 %	3
prob. false behaviors	0.0124	0.0180	

Table 1: M/G/2/3/3 queue: Errors due to the approximation of class domains.

Approximation of state density functions affects a higher number of classes, since it may be applied both at the end of time advancement, when the state density function accepts a piecewise representation over a partition in polyhedral subdomains, and at the end of disabling, when the state density function accepts a piecewise representation over a partition in DBM subdomains. As shown in Table 2, in our example, state density functions are approximated after the steps of time advancement and disabling in 36 and 30 classes, respectively, yielding a d-norm distance with average value of $7.7 \cdot 10^{-7}$ and $4.487 \cdot 10^{-4}$, respectively, and maximum value of $1.62 \cdot 10^{-5}$ and $1.9278 \cdot 10^{-3}$, respectively.

Fig. 3 compares steady state probabilities of reachable marking obtained by approximated analysis with simulation results obtained after 200000 firings. The overall fit together with average and maximum error, equal to 0.00252 and 0.00510, respectively, evidence that analysis results agree with simulation results.

Metrics on densities	average	max	# classes
d-norm distance (time advancement)	0.0000077	0.0000162	36
d-norm distance (disabling)	0.0004487	0.0019278	30

Table 2: M/G/2/3/3 queue: Errors due to the approximation of state density functions.

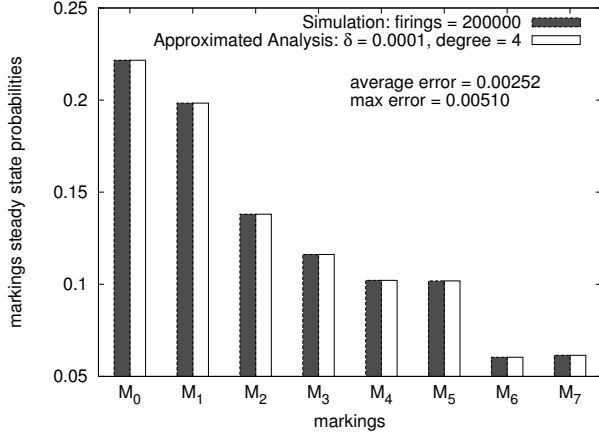


Figure 3: M/G/2/3/3 queue: Measures of steady state marking probabilities obtained by simulation (with 200000 firings) and approximated analysis (with threshold $\delta = 0.0001$ and Bernstein degree 4).

4.2 A complex example

Fig. 4 depicts a variant of the Three-Tasks model introduced in [20], called *Three-Synchronizing-Tasks*, which represents three concurrent tasks that perform three computations at each activation and share three mutually exclusive resources. The three tasks are made by transitions $t1x$, $t2x$, and $t3x$, respectively: transitions $tx0$ model tasks arrival; transitions $tx2$, $tx3$ and $tx5$ represent three computations performed at each task activation; transitions $tx1$ and $tx4$ model the acquisition of a mutually exclusive resource which is necessary to perform the subsequent computations $tx2$ and $tx5$, respectively. In particular, these mutually exclusive resources are represented by places $R12$, $R23$, and $R13$. Times-to-fire of transitions $tx0$ are exponentially distributed over $[0, \infty]$ with rate 0.003; times-to-fire of transitions $tx1$ and $tx4$ are expolynomially distributed over $[0, 100]$ according to $f(x) = k x e^{-0.003x}$; times-to-fire of transitions $tx2$ and $tx5$ are uniformly distributed over $[200, 400]$; and, time-to-fire of transitions $tx3$ are uniformly distributed over $[100, 200]$. The first and the third task require resource $r1$ with priority level 1 and 2, respectively; the second task requires resource $r2$ with priority level 1; all tasks require resource $r3$ whenever they need to access one of the three mutually exclusive resources $R12$, $R23$, and $R13$, raising their own priority to the highest priority level of any task that may access the resource.

Approximated analysis enumerates 537 stochastic classes, covered by 180 markings, in nearly 2 minutes. Approximation of domains affects 14 classes, as shown in Table 3: also in this case, although the estimated average and maximum portion of false behaviors in any of these classes is a notable quantity, equal to 26.14 % and

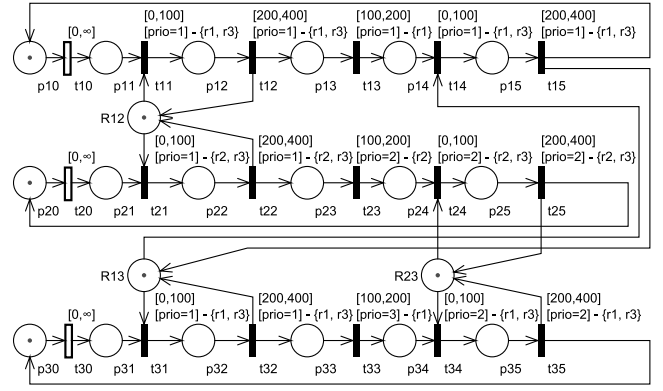


Figure 4: Three-Synchronizing-Tasks model: An spTPN model representing three recurrent tasks sharing three resources.

47.00 %, respectively, their probability is quite low and, thus, does not impair performance measures on the model.

Metrics on domains	average	max	# classes
% false behaviors	26.14 %	47.00 %	14
prob. false behaviors	0.0240	0.0512	

Table 3: Three-Synchronizing-Tasks model: Errors due to the approximation of class domains.

As illustrated in Table 4, approximation of state density functions affects a consistent number of classes (i.e., 201 after time advancement and 149 after disabling), yielding average and maximum errors under acceptable limits ($5.7 \cdot 10^{-6}$ and $3.09 \cdot 10^{-5}$, $3.376 \cdot 10^{-4}$ and $1.9885 \cdot 10^{-3}$, for time advancement and disabling steps, respectively).

Metrics on densities	average	max	# classes
d-norm distance (time advancement)	0.0000057	0.0000309	201
d-norm distance (disabling)	0.0003376	0.0019885	149

Table 4: Three-Synchronizing-Tasks model: Errors due to the approximation of state density functions.

Fig. 5 reports results obtained on steady state marking probabilities, evidencing acceptable errors with respect to simulation results, with average and maximum error equal to 0.00023 and 0.00211, respectively.

Experiments were repeated with different values of the approximation threshold δ (i.e., 0.001 and 0.00001) and the Bernstein degree (i.e., 3 and 5). Approximated analysis is substantially insensitive to the threshold and, instead, yields more accurate quantitative measures as the Bernstein degree increases. In particular, with degree 3: average and maximum probability of false behaviors are equal to 0.0329 and 0.0649, respectively; average and maximum d-norm distance after time-advancement are equal to 0.0000066 and 0.0000343, respectively; average and maximum d-norm distance after disabling are equal to 0.0004057 and 0.0019314, respectively. With degree 5: average and maximum probability of

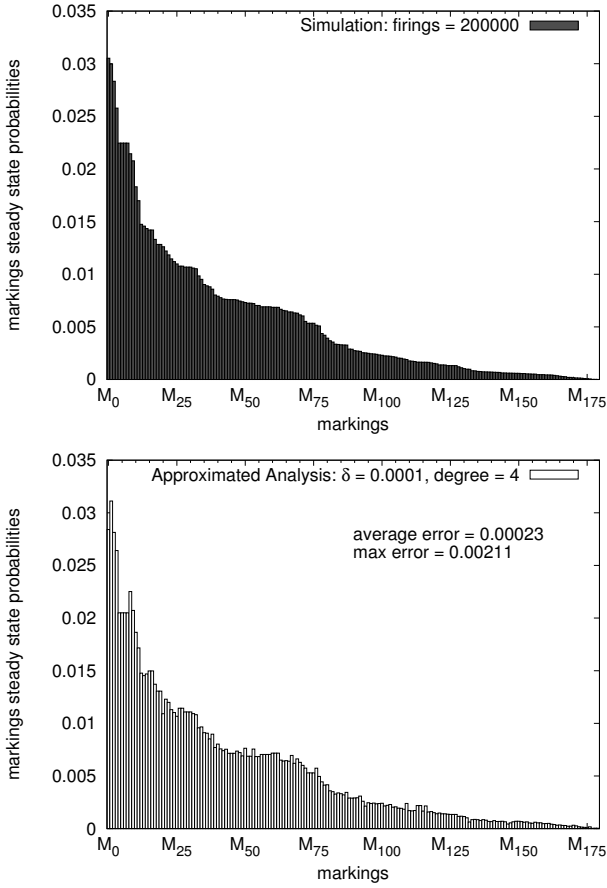


Figure 5: Three-Tasks model: Measures of steady state marking probabilities obtained by simulation (with 200000 firings) and approximated analysis (with threshold $\delta = 0.0001$ and Bernstein degree 4).

false behaviors are equal to 0.0179 and 0.0413, respectively; average and maximum d-norm distance after time-advancement are equal to 0.0000050 and 0.0000286, respectively; average and maximum d-norm distance after disabling are equal to 0.00029023 and 0.0013614, respectively.

5. CONCLUSIONS

We proposed an analytical approach for quantitative evaluation of systems with multiple concurrently enabled GEN timers running under fixed-priority preemptive scheduling. The approach extends the theory of stochastic state classes so as to deal with models that combine Preemptive Repeat Different (PRD) and Preemptive Resume (PRs) policies. The expressive extension exacerbates the complexity of analysis, as it requires to deal with timer vectors distributed in piecewise form over linear convex polyhedra. To reduce complexity, the approach approximates exact distributions with global Bernstein approximants supported over DBM zones. Computational experience illustrates application on a relatively complex model that combines usual patterns of real-time concurrency and non-Markovian temporal parameters, showing that the approach attains a significant reduction of complexity while suffering a limited impact on the accuracy of quantitative measures.

The representation of preemptive behavior is a relevant issue also in non-deterministic analysis, where approximation seems to be the only viable approach to manage this class of models as well. In stochastic analysis, the problem becomes even more complex both in the theory and in the practical implementation. However, quantitative measures also provide a notable leverage to restrain the impact of approximation, since false behaviors are associated with a measure of probability. Experimental results show that while the enlargement of domains can be significant, its impact on performance measures is still limited.

Appendix: Theorem Proof

Theorem 3.1. Integration bounds $E_{\underline{x}_G}^{\beta_0}(\underline{x}_G^c)$ and $L_{\underline{x}_G}^{\beta_0}(\underline{x}_G^c)$ are piecewise continuous functions with analytic representation over each element of a partition of $D_{\underline{x}_G}^c$ in polyhedral subdomains R_{ixjy} , with $i, j \in [0, I-1]$ and $x, y \in [I, G-1]$. Within each subdomain R_{ixjy} , the two integration bounds are either constants, or mono-variate linear functions of a single component of the vector \underline{x}_G^c , or bivariate linear functions of two components of the vector \underline{x}_G^c . In particular, if $\underline{x}_G^c \in R_{ixjy}$, then:

$$E_{\underline{x}_G}^{\beta_0}(\underline{x}_G^c) = \begin{cases} -B_{*\beta_0} & \text{iff } i = 0, x = I \\ -B_{*\beta_i} - x_{\beta_i} & \text{iff } i > 0, x = I \\ -B_{\beta_x\beta_0} + x_{\beta_x} & \text{iff } i = 0, x > I \\ -B_{\beta_x\beta_i} + x_{\beta_x} - x_{\beta_i} & \text{iff } i > 0, x > I \end{cases}$$

$$L_{\underline{x}_G}^{\beta_0}(\underline{x}_G^c) = \begin{cases} B_{\beta_0*} & \text{iff } j = 0, y = I \\ B_{\beta_j*} - x_{\beta_j} & \text{iff } j > 0, y = I \\ B_{\beta_0\beta_y} + x_{\beta_y} & \text{iff } j = 0, y > I \\ B_{\beta_j\beta_y} + x_{\beta_y} - x_{\beta_j} & \text{iff } j > 0, y > I \end{cases} \quad (26)$$

PROOF. Let $\{\bar{R}_{ix}\}$ with $i \in [0, I-1]$ and $x \in [I, G-1]$ be subsets of $D_{\underline{x}_G}^c$ defined as:

$$\begin{aligned} \bar{R}_{0I} &\stackrel{def}{=} D_{\underline{x}_G}^c \cap \{ \underline{x} \mid x_{\beta_i} \leq B_{*\beta_i} + B_{*\beta_0} \forall i \in [1, I-1] \\ &\wedge -x_{\beta_x} \leq B_{*\beta_0} - B_{\beta_x\beta_0} \forall x \in [I, G-1] \\ &\wedge x_{\beta_i} - x_{\beta_x} \leq B_{*\beta_0} + B_{\beta_x\beta_i} \\ &\forall i \in [1, I-1], \forall x \in [I, G-1] \}, \end{aligned}$$

$$\begin{aligned} \bar{R}_{iI} &\stackrel{def}{=} D_{\underline{x}_G}^c \cap \{ \underline{x} \mid -x_{\beta_i} \leq B_{*\beta_i} - B_{*\beta_0} \\ &\wedge x_{\beta_j} - x_{\beta_i} \leq B_{*\beta_i} - B_{*\beta_j} \forall j \in [1, I-1] \cdot j \neq i \\ &\wedge -x_{\beta_i} - x_{\beta_x} \leq B_{*\beta_i} - B_{\beta_x\beta_0} \forall x \in [I, G-1] \\ &\wedge x_{\beta_i} + x_{\beta_j} - x_{\beta_x} \leq B_{*\beta_i} - B_{\beta_x\beta_j} \\ &\forall j \in [1, I-1] \cdot j \neq i, \forall x \in [I, G-1] \}, \end{aligned}$$

$$\begin{aligned} \bar{R}_{0x} &\stackrel{def}{=} D_{\underline{x}_G}^c \cap \{ \underline{x} \mid x_{\beta_x} \leq B_{\beta_x\beta_0} - B_{*\beta_0} \\ &\wedge x_{\beta_x} + x_{\beta_i} \leq B_{\beta_x\beta_0} - B_{*\beta_i} \forall i \in [1, G-1] \cdot i \neq j \\ &\wedge x_{\beta_x} - x_{\beta_y} \leq B_{\beta_x\beta_0} - B_{\beta_y\beta_0} \forall y \in [I, G-1] \cdot y \neq x \\ &\wedge x_{\beta_x} - x_{\beta_y} + x_{\beta_i} \leq B_{\beta_x\beta_0} - B_{\beta_y\beta_i} \\ &\forall i \in [1, H-1], \forall y \in [I, G-1] \cdot y \neq x \}, \end{aligned}$$

$$\begin{aligned} \bar{R}_{ix} &\stackrel{def}{=} D_{\underline{x}_G}^c \cap \{ \underline{x} \mid x_{\beta_x} - x_{\beta_i} \leq B_{\beta_x\beta_i} - B_{*\beta_0} \\ &\wedge x_{\beta_x} - x_{\beta_i} + x_{\beta_j} \leq B_{\beta_x\beta_i} - B_{*\beta_j} \\ &\forall j \in [I, G-1] \cdot j \neq i \\ &\wedge x_{\beta_x} - x_{\beta_i} - x_{\beta_y} \leq B_{\beta_x\beta_i} - B_{\beta_y\beta_0} \\ &\forall y \in [I, G-1] \cdot y \neq x \\ &\wedge x_{\beta_x} - x_{\beta_i} - x_{\beta_y} + x_{\beta_j} \leq B_{\beta_x\beta_i} - B_{\beta_y\beta_j} \\ &\forall j \in [1, I-1] \cdot j \neq i, \forall y \in [I, G-1] \cdot y \neq x \}. \end{aligned} \quad (27)$$

According to Eq.(20-a):

$$E_{\mathcal{I}_G}^{\beta_0}(\underline{x}) = \begin{cases} -B_{*\beta_0} & \text{iff } \underline{x} \in \bar{R}_{0I} \\ -B_{*\beta_i} - x_{\beta_i} & \text{iff } \underline{x} \in \bar{R}_{iI} \\ -B_{\beta_x\beta_0} + x_{\beta_x} & \text{iff } \underline{x} \in \bar{R}_{0x} \\ -B_{\beta_x\beta_i} + x_{\beta_x} - x_{\beta_i} & \text{iff } \underline{x} \in \bar{R}_{ix}. \end{cases} \quad (28)$$

Regions \bar{R}_{ixjy} are convex polyhedral zones, being the restriction of a convex polyhedron through linear constraints that may not be in DBM form; they are disjoint by construction; their union is $D_{\mathcal{I}_G}$; and, within each of them, $E_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ is either constant (i.e., $-B_{*\beta_0}$, in \bar{R}_{0I}), or a monovariate linear function (i.e., $-B_{*\beta_i} - x_{\beta_i}$ or $x_{\beta_x} - B_{\beta_x\beta_0}$, in \bar{R}_{iI} or \bar{R}_{0x} , respectively), or a bivariate linear function (i.e., $x_{\beta_x} - B_{\beta_x\beta_i} - x_{\beta_i}$, in \bar{R}_{ix}). In a similar manner, Eq.(20-b) partitions $D_{\mathcal{I}_G}$ into a finite set of polyhedral subdomains $\{\bar{R}_{jy}\}$ with $j \in [0, I-1]$ and $y \in [I, G-1]$, within each of which $L_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ accepts analytic representation. The intersection of the two partitions defines a set of polyhedral subdomains R_{ixjy} with $i, j \in [0, I-1]$ and $x, y \in [I, G-1]$, within each of which both $E_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ and $L_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ have analytic representation.

Given a point \hat{x} at the border between two regions, say \bar{R}_{ix} and \bar{R}_{jy} , Eqs.(27) and (28) imply that:

$$\begin{aligned} \hat{x}_{\beta_x} - \hat{x}_{\beta_i} - B_{\beta_x\beta_i} &= \hat{x}_{\beta_y} - \hat{x}_{\beta_j} - B_{\beta_y\beta_j} \\ \lim_{\underline{x} \rightarrow \hat{x}, \underline{x} \in \bar{R}_{ix}} E_{\mathcal{I}_G}^{\beta_0}(\underline{x}) &= \lim_{\underline{x} \rightarrow \hat{x}} \hat{x}_{\beta_x} - B_{\beta_x\beta_i} - \hat{x}_{\beta_i} \\ &= \hat{x}_{\beta_y} - B_{\beta_y\beta_j} - \hat{x}_{\beta_j} \quad (29) \\ \lim_{\underline{x} \rightarrow \hat{x}, \underline{x} \in \bar{R}_{jy}} E_{\mathcal{I}_G}^{\beta_0}(\underline{x}) &= \lim_{\underline{x} \rightarrow \hat{x}} \hat{x}_{\beta_y} - B_{\beta_y\beta_j} - \hat{x}_{\beta_j} \\ &= \hat{x}_{\beta_x} - B_{\beta_x\beta_i} - \hat{x}_{\beta_i} \end{aligned}$$

According to this, $E_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ is continuous along subdomains borders and, since it is analytic within each subdomain R_{ixjy} , it is a globally continuous function over $D_{\mathcal{I}_G}$. The same steps can be applied to demonstrate that $L_{\mathcal{I}_G}^{\beta_0}(\underline{x})$ is a globally continuous function. \square

6. REFERENCES

- [1] A. Zimmermann and J. Freiheit and R. German and G. Hommel. Petri Net Modelling and Performability Evaluation with TimeNET 3.0. *11th Int. Conf. on Modelling Techniques and Tools for Comp. Perf. Evaluation (TOOLS'2000)*, pages 188–202, 2000.
- [2] G. Ciardo and R. German and C. Lindemann. A characterization of the stochastic process underlying a stochastic Petri net. *IEEE Trans. On SW Eng.*, 20(7):506–515, 1994.
- [3] A. Bobbio and A. Puliafito and M. Telek. A Modelling Framework to implement preemption policies in non-Markovian SPNs. *IEEE Trans. On SW Eng.*, 26(1), January 2000.
- [4] A. Horvath and A. Puliafito and M. Scarpa and M. Telek. Analysis and Evaluation of NonMarkovian Stochastic Petri Nets. In *Proc. of the 11th Int. Conf. Computer Performance Evaluation (TOOLS)*, pages 171–187, 2000.
- [5] Alan Burns and Sanjoy Baruah. Sustainability in Real-time Scheduling. *Journal of Computing Science and Engineering*, 2(1):74–97, March 2008.
- [6] Andrea Bobbio and Miklos Telek. Transient Analysis of a Preemptive Resume M/D/1/2/2 through Petri Nets. *Periodica Politechnica*, 41:123–146, September 1997.
- [7] B. Berthomieu and M. Diaz. Modeling and Verification of Time Dependent Systems Using Time Petri Nets. *IEEE Trans. on SW Eng.*, 17(3):259–273, March 1991.
- [8] L. Carnevali, L. Ridi, and E. Vicario. Putting preemptive Time Petri Nets to work in a V-Model SW life cycle. *IEEE Trans. on SW Eng.*, *accepted for publication*.
- [9] E. Teruel and G. Franceschinis and M. De Pierro. Well defined Generalized Stochastic Petri Nets: A Net Level Method To Specify Priorities. *IEEE Trans. On Soft. Eng.*, 29(11):465 – 478, 2003.
- [10] E. Vicario. Static Analysis and Dynamic Steering of Time Dependent Systems Using Time Petri Nets. *IEEE Trans. on SW Eng.*, 27(1):728–748, August 2001.
- [11] E. Vicario and L. Sassoli and L. Carnevali. Using Stochastic State Classes in Quantitative Evaluation of Dense-Time Reactive Systems. *IEEE Trans. on SW Eng.*, 35(5):703–719, 2009.
- [12] F. Cassez and K. G. Larsen. *The Impressive Power of Stopwatches*, volume 1877. LNCS, August, 2000.
- [13] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1998.
- [14] G. Bucci and A. Fedeli and L. Sassoli and E. Vicario. Timed State Space Analysis of Real Time Preemptive Systems. *IEEE Trans. SW Eng.*, 30(2):97–111, Feb. 2004.
- [15] G. Bucci and L. Carnevali and L. Ridi and E. Vicario. Oris: a Tool for Modeling, Verification and Evaluation of Real-Time Systems. *Int. Journal of Software Tools for Technology Transfer*, 12(5):391 – 403, 2010.
- [16] G. Bucci and L. Sassoli and E. Vicario. Correctness verification and performance analysis of real time systems using stochastic preemptive Time Petri Nets. *IEEE Trans. on SW Eng.*, 31(11):913–927, November 2005.
- [17] G. Ciardo and A. Blakemore and P.F. Chimento and J.K. Muppala and K.S. Trivedi. Automated generation and analysis of Markov reward models using Stochastic Reward Nets. *Linear Algebra, Markov Chains, and Queueing Models, IMA Volumes in Mathematics and its Applications*, 48:145–191, 1993.
- [18] G.G. Lorentz. *Bernstein Polynomials*. University of Toronto Press, 1953.
- [19] L. Carnevali and L. Grassi and E. Vicario. A Tailored V-Model Exploiting the Theory of Preemptive Time Petri Nets. In *Proc. of the Ada-Europe Int. Conf. on Reliable SW Tech.*, pages 87–100. Springer-Verlag, 2008.
- [20] L. Carnevali and L. Grassi and E. Vicario. State-Density Functions over DBM Domains in the Analysis of Non-Markovian Models. *IEEE Trans. on SW Eng.*, 35(2):178–194, 2009.
- [21] D. Lime and O. H. Roux. Formal verification of real-time systems with preemptive scheduling. *Real-Time Syst.*, 41(2):118–151, 2009.
- [22] Olivier H. Roux and Didier Lime. Time Petri nets with inhibitor hyperarcs: formal semantics and state-space computation. *25th Int. Conf. on Theory and Application of Petri Nets*, 3099:371–390, 2004.
- [23] Tomas Sauer. Multivariate Bernstein polynomials and convexity. *Computer Aided Geometric Design*, 8(6):465 – 478, 1991.