

©2005 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Introducing Probability within State Class Analysis of Dense-Time-Dependent Systems

G.Bucci, R.Piovosi, L.Sassoli, E.Vicario

Abstract

Several techniques have been proposed for symbolic enumeration and analysis of the state space of reactive systems with non-deterministic temporal parameters taking values within a dense domain. In a large part of these techniques, the state space is covered by collecting states within equivalence classes each comprised of a discrete logical location and a dense variety of clock valuations encoded as a Difference Bounds Matrix (DBM). The reachability relation among such classes enables qualitative verification of properties pertaining the ordering of events along critical runs and the satisfaction of stimulus/response deadlines. However, up to now, no results have been proposed which extend state class enumeration so as to combine the verification of the possibility of critical behaviors with a quantitative evaluation of their probability.

In this paper, we extend the concept of equivalence classes based on DBM encoding with a density function which provides a measure for the probability associated with individual states collected in the class itself. To this end, we extend the formalism of Time Petri Nets by associating the static firing interval of each transition with a probability density function. We then expound how this probabilistic information determines a probability for the states collected within a class and how this probability evolves in the enumeration of the reachability relation among state classes. This opens the way to characterizing the possibility of critical behaviors with a quantitative measure of probability.

1 Introduction

Development of reactive and time-dependent systems jointly addresses requirements pertaining ordered sequencing of events, stimulus-response timeliness, and efficient resource usage [21] [31] [29]. Despite this demand raised by the application domain, modeling and analysis techniques for correctness verification and for performance/dependability evaluation have been separately ad-

ressed in different timed variants of Petri Nets [13].

On the one hand, in the context of performance and dependability evaluation, stochastic Petri Nets associate timed transitions with a stochastic delay characterized through an exponential density function [26] [1]. This enables Markovian analysis and permits automated derivation of effective performance and dependability indexes [17]. As a major drawback, the unbounded support of the exponential distribution does not permit the modeler to represent implicit precedences induced by finite timing constraints (e.g. time-outs).

Several extensions of stochastic Petri Nets have been developed to encompass bounded delays and to partially overcome the limits of exponential timing[9][20]. However, the application of these techniques imposes various restrictions which exclude models allowing multiple concurrent non-exponential clocks [2] [15] [16] [11], or models where timing constraints are essential to keep the set of reachable markings finite [27] [10].

On the other hand, in the context of correctness verification of real time systems, a number of analysis techniques have been proposed for models such as Timed Automata and Time Petri Nets which include non-deterministic temporal parameters taking values within (possibly finite) dense intervals [4][6] [18][7][32]. For this kind of models, the timed state space is covered through the enumeration of a discrete reachability relation among state classes, each comprised of a discrete logical location and a time domain collecting a dense variety of timings. In particular, a wide literature has been developed upon state classes where time domains are encoded as difference bounds matrixes (DBM) [3] [8] [7] [19] [5]. Enumeration and analysis of the reachability relation among such state classes opens the way to the solution of a number of relevant problems, such as the reachability of a given logical location, the feasibility of a run satisfying given constraints on the logical sequencing of events and on their quantitative timing, the evaluation of a tight bound on the minimum and maximum time that can elapse between any two events along a symbolic run [7] [6] [24] [32].

However, these techniques do not permit the characterization of feasible behaviors with a measure of probability,

which is an essential step towards dependability and performance evaluation. To overcome this limitation, probabilistic extensions of timed automata have been proposed to enable model checking for probabilistic real time systems with continuous random delays, allowing the verification of quantitative probability bounds [23] [12][25].

In this paper, we address the problem of deriving a density function which characterizes the probability of individual timings comprised within the boundaries of a time domain in DBM form. To this end, we extend the formalism of Time Petri Nets by associating the static firing interval of each transition with a (dense) probability density function. We then expound how this probabilistic information induces a measure of probability for individual states collected in a class and how this probability evolves in the enumeration of the reachability relation among state classes.

The rest of the paper is organized in four sections. Time Petri Nets extended with stochastic time intervals are defined in Sect.2. In Sects.3 and 4, we extend the concept of state class with a density function capturing the probability of individual states in the class, we present a method for the derivation of the successors of a stochastic state class, and we discuss the application of this derivation within an enumerative semi-algorithm. Conclusions are drawn in Sect.5.

2 Time Petri Nets with stochastic firing intervals

A Stochastic Time Petri Net (sTPN) is a tuple

$$sTPN = \langle P; T; A^-; A^+; M; A; D \rangle \quad (1)$$

- The first seven members comprise the basic model of Time Petri Nets: P is a set of *places*; T a set of *transitions*; A^- and A^+ are sets of preconditions and postconditions connecting places to transitions and vice versa, respectively:

$$\begin{aligned} A^- &\subseteq P \times T \\ A^+ &\subseteq T \times P \end{aligned} \quad (2)$$

A place p is said to be an *input* or an *output* place for a transition t if there exists a precondition or a postcondition from p to t or vice versa, (i.e. if $\langle p, t \rangle \in A^-$ or $\langle t, p \rangle \in A^+$), respectively. M (the initial marking) associates each place with a non-negative number of tokens:

$$M : P \rightarrow \mathbb{N} \cup \{0\} \quad (3)$$

P , T , A^- , and A^+ comprise a bipartite graph, P and T being disjoint classes of nodes, and A^- and

A^+ being relations between them. This graph is represented graphically by drawing places as circles, transitions as bars, and preconditions and postconditions as directed arcs; the tokens of the initial marking are represented as dots inside places.

A^- is a set of inhibitor arcs connecting places to transitions:

$$A^- \subseteq P \times T \quad (4)$$

inhibitor arcs are represented graphically as dot-terminated arcs.

- \mathcal{D} associates each transition t with a static probability density function \mathcal{D}_t ($\mathcal{D} : T \rightarrow \Omega$, where Ω denotes the space of probability density functions). For each transition t , the support of \mathcal{D}_t represents a *static firing interval*, which is made up of an *earliest* and a (possibly infinite) *latest firing time*, called $EFT^s(t)$ and $LFT^s(t)$, respectively.

The *state* of a sTPN is a pair $s = \langle M, \tau \rangle$, where M is the *marking* and τ associates each transition with a possibly infinite *time to fire* value ($\tau : T \rightarrow \mathbb{R}^+ \cup \{\infty\}$). The state evolves according to a transition rule made up of two clauses of *firability* and *firing*.

Firability: A transition t_o is *enabled* if each of its input places contains at least one token and none of its inhibiting places contains any token. A transition t_o is *firable* if its time to fire $\tau(t_o)$ is not higher than the time to fire of any other progressing transition.

Firing: When a transition t_o fires, the state $s = \langle M, \tau \rangle$ is replaced by a new state $s' = \langle M', \tau' \rangle$. The marking M' is derived from M by removing a token from each input place of t_o , and by adding a token to each output place of t_o :

$$\begin{aligned} M_{tmp}(p) &= M(p) - 1 \quad \forall p. \langle p, t_o \rangle \in A^- \\ M'(p) &= M_{tmp}(p) + 1 \quad \forall p. \langle t_o, p \rangle \in A^+ \end{aligned} \quad (5)$$

Transitions that are enabled both by the temporary marking M_{tmp} and by the final marking M' are said *persistent*, while those that are enabled by M' but not by M_{tmp} are said *newly enabled*. If t_o is still enabled after its own firing, it is always regarded as newly enabled.

The time to fire τ' of any transition enabled by the new marking M' is computed in a different manner for newly enabled transitions and for persistent transitions:

i) for transition t_a which is newly enabled after the firing of t_o , the time to fire takes a nondeterministic value sampled in the static firing interval, according to the static probability density function $f_{t_a}(\cdot)$:

$$EFT^s(t_a) \leq \tau'(t_a) \leq LFT^s(t_a) \quad (6)$$

ii) for any transition t_i which is persistent after the firing of t_o , the time to fire is reduced by the time elapsed in the

previous state. This is equal to the time to fire of t_o as it was measured at the entrance in the previous state:

$$\tau'(t_i) = \tau(t_i) - \tau(t_o) \quad (7)$$

3 Augmenting State Class With Probability

3.1 States, State Classes and Stochastic State Classes

In the firing clause of sTPNs, a newly enabled transition may take any real value within its static firing interval, and each value can lead to a different state, thus resulting in a dense variety of possible successors. To obtain a discrete representation of the state space, the reachability relation between states is conveniently replaced through some reachability relation between *state classes*, each made up by a dense variety of states with the same marking m but with different timings comprised within a *firing domain* D_m [19][32]:

$$\text{State class} = \langle m, D_m \rangle \quad (8)$$

The encoding of the firing domain D_m jointly depends on the way in which transition timers are made to advance in the firing clause and on the semantics of the reachability relation established among state classes. Most works (and this among them) on the analysis of densely timed models are based on the EA reachability relation [28]:

Definition 3.1 class S^c is a successor of class S^p through t_o (which is also written as $S^p \xrightarrow{t_o} S^c$) if and only if S^c contains all and only the states that are reachable from some states collected in S^p through some feasible firing of t_o .

Under this reachability relation, the firing domain of state classes of a TPN model can be represented as the set D of solutions of a set of linear inequalities in the form of a Difference Bounds Matrix (DBM) [19]:

$$D = \left\{ \begin{array}{l} \tau(t_i) - \tau(t_j) \leq b_{ij} \\ \forall t_i, t_j \in T(m) \cup \{t_*\} \quad t_i \neq t_j \end{array} \right. \quad (9)$$

where $T(m)$ denotes the set of transitions enabled by m , $\tau(t_i)$ denotes the time to fire of transition t_i , the fictitious unknown variable $\tau(t_*) = 0$ serves to keep all the inequalities in the same difference form, and $b_{ij} \in \mathbb{R} \cup \{+\infty\}$ are the coefficients which define the boundaries of a class. The DBM form has a *normal* representation which can be computed as the solution of an *all shortest path* problem, and which supports efficient detection and derivation of successor classes, in time $O(N)$ and $O(N^2)$ respectively, with respect to the number of enabled transitions [32].

The DBM representation can be applied to encode the range of feasible timings of an sTPN, as the support of feasible timings of this model evolves with the same semantics as a TPN. However, this encoding does not exploit

the stochastic information which is introduced in sTPNs to characterize the probability of different determinations of temporal parameters. To overcome the limitation, we introduce a concept of *stochastic state class* which extends a state class $\langle m, D \rangle$ with a joint probability function $f_{\vec{\tau}}(\cdot)$ characterizing the distribution of the vector $\vec{\tau} = \langle \tau(t_0), \tau(t_1), \dots, \tau(t_n) \rangle$ of times to fire of transitions enabled by m within the limits of the firing domain D :

$$\text{Stochastic state class} = \langle m, D, f_{\vec{\tau}}(\cdot) \rangle \quad (10)$$

Each determination of $\vec{\tau}$ uniquely identifies a state in class S and vice versa. According to this, $f_{\vec{\tau}}(\cdot)$ takes the meaning of a density function for the probability of the states in S , for which we call it *state probability density function*.

With this perspective, we extend the reachability relation among state classes as follows:

Definition 3.2 given two stochastic state classes $\Sigma^p = \langle m^p, D^p, f_{\vec{\tau}^p}(\cdot) \rangle$ and $\Sigma^c = \langle m^c, D^c, f_{\vec{\tau}^c}(\cdot) \rangle$, we say that Σ^c is a successor of Σ^p through t_o with probability μ , and we write $\Sigma^p \xrightarrow{t_o} \Sigma^c$, iff the following property holds: if the marking of the net is m^p and the vector of times to fire of transitions enabled by m^p is a random variable $\vec{\tau}^p$ distributed within the boundaries of D^p according to $f_{\vec{\tau}^p}(\cdot)$, then t_o is a possible firing, which occurs with probability μ and which leads to a new marking m^c and a new vector of times to fire distributed within the boundaries of D^c according to $f_{\vec{\tau}^c}(\cdot)$.

In the following, we develop the steps for the enumeration of this reachability relation, i.e. the detection of successors, the calculus of their probability, and the derivation of successor state-probability density functions.

3.2 Successors detection and calculus of their probability

A transition t_o is an outgoing event from the stochastic class $\Sigma^p = \langle m^p, D^p, f_{\vec{\tau}^p}(\cdot) \rangle$ iff t_o is enabled by the marking m^p and the firing domain D^p accepts solutions in which the firing time $\tau(t_o)$ of transition t_o is not greater than that of any other enabled transition. This occurs iff the following *restricted firing domain* $D_{t_o}^p$ accepts a non-empty set of solutions:

$$D_{t_o}^p = \left\{ \begin{array}{l} \tau(t_i) - \tau(t_j) \leq b_{ij} \\ \tau(t_o) - \tau(t_j) \leq \min\{0, b_{oj}\} \\ \forall t_i, t_j \in T(m^p) \cup \{t_*\} \quad t_i \neq t_j \end{array} \right. \quad (11)$$

If t_o is a possible outgoing event, its probability μ is derived by integrating the state density function $f_{\vec{\tau}^p}(\cdot)$ over the restricted firing domain $D_{t_o}^p$:

$$\begin{aligned} \mu &= \text{Prob}\{t_o \text{ fires first}\} = \\ &= \text{Prob}\{\vec{\tau} \in D_{t_o}^p\} = \int_{D_{t_o}^p} f_{\vec{\tau}^p}(\vec{x}) d\vec{x} \end{aligned} \quad (12)$$

3.3 Derivation of successor state-probability density functions

In the computation of the class $\Sigma^c = \langle m^c, D^c, f_{\vec{\tau}^c}(\cdot) \rangle$ reached from $\Sigma^p = \langle m^p, D^p, f_{\vec{\tau}^p}(\cdot) \rangle$ through an outgoing event t_o (i.e. $\Sigma^p \xrightarrow{t_o} \Sigma^c$), the new marking m^c is derived by moving tokens according to the execution rule of transitions, and the firing domain D^c is derived so as to reflect the evolution of times to fire. Details of the derivation are reported in [32]. For the present treatment, it is sufficient to resume the steps of the derivation as follows:

1. the vector of times to fire $\vec{\tau}^p = \langle \tau(t_o), \tau(t_1), \dots, \tau(t_n) \rangle$ of the transitions enabled in S^p is replaced with the vector $\vec{\tau}' = \langle \tau'(t_o), \tau'(t_1), \dots, \tau'(t_n) \rangle$ where each unknown value $\tau'(t_i)$ is obtained by restricting $\tau(t_i)$ with the constraint $\tau(t_i) \geq \tau(t_o)$ so as to capture the condition for t_o to be the firing transition;
2. $\vec{\tau}'$ is replaced through the vector $\vec{\tau}'' = \langle \tau''(t_o), \tau''(t_1), \dots, \tau''(t_n) \rangle = \langle \tau'(t_o), \tau'(t_1) - \tau'(t_o), \dots, \tau'(t_n) - \tau'(t_o) \rangle$ so as to reflect the reduction of times to fire during the permanence in the parent class Σ^p ;
3. times to fire of enabled transitions at the firing of t_o are obtained by eliminating $\tau(t_o)$ from $\vec{\tau}''$ through a projection operation which yields a new vector $\vec{\tau}''' = \langle \tau'''(t_1), \dots, \tau'''(t_n) \rangle$;
4. the vector $\vec{\tau}''''$ of times to fire in the child class Σ^c is finally obtained by removing through a projection the times to fire of transitions that are not persistent after the firing of t_o and by adding the times to fire of newly enabled transitions, each constrained within its own static firing interval.

Derivation of the probability density function within the boundaries of the firing domain of the child class Σ^c can be organized along the same four steps, extending the derivation of inequalities with a stochastic characterization of their solution space.

1. We regard $\vec{\tau} = \langle \tau(t_o), \tau(t_1), \dots, \tau(t_n) \rangle$ as a stochastic array variable, and $\vec{\tau}' = \langle \tau'(t_o), \tau'(t_1), \dots, \tau'(t_n) \rangle$ as the variable obtained by conditioning $\vec{\tau}$ through the assumption that t_o will fire first, i.e. that $\tau(t_o) \leq \tau(t_i)$ for any enabled transition t_i in Σ^p :

$$\begin{aligned} \vec{\tau}' &= \langle \tau'(t_o), \dots, \tau'(t_n) \rangle \\ \tau'(t_i) &= \tau(t_i) \mid \tau(t_o) \leq \tau(t_i) \quad \forall i = 1, n \end{aligned} \quad (13)$$

The joint density function of $\vec{\tau}'$ can be expressed through Bayes Theorem as:

$$f_{\vec{\tau}'}(\tau'_o, \dots, \tau'_n) = \begin{cases} \frac{f_{\vec{\tau}}(\tau'_o, \dots, \tau'_n)}{\int_{D_{t_o}^p} f_{\vec{\tau}}(\tau_o, \dots, \tau_n) d\tau_o \dots d\tau_n} & \text{if } \tau'_o, \dots, \tau'_n \in D_{t_o}^p \\ 0 & \text{if } \tau'_o, \tau'_1, \dots, \tau'_n \notin D_{t_o}^p \end{cases} \quad (14)$$

2. The stochastic array variable $\vec{\tau}''$ is obtained by replacing each stochastic variable $\tau'(t_i)$ with $i > 0$ through the difference $\tau'(t_i) - \tau'(t_o)$:

$$\tau''(t_i) = \begin{cases} \tau'(t_i) - \tau'(t_o) & \forall i = 1 \dots, n \\ \tau'(t_o) & \text{for } i = 0 \end{cases} \quad (15)$$

The joint density function $f_{\vec{\tau}''}$ of the variable $\vec{\tau}''$ can be expressed as:

$$f_{\vec{\tau}''}(\tau''_o, \tau''_1, \dots, \tau''_n) = f_{\vec{\tau}'}(\tau''_o, \tau''_1 + \tau''_o, \dots, \tau''_n + \tau''_o) \quad (16)$$

3. The stochastic variable $\vec{\tau}'''$ is derived from $\vec{\tau}''$ through a projection eliminating the variable $\tau''(t_o)$:

$$\vec{\tau}''' = \langle \tau''(t_1), \tau''(t_2), \dots, \tau''(t_n) \rangle \quad (17)$$

The joint density function $f_{\vec{\tau}'''}$ is thus be obtained by integrating the density function $f_{\vec{\tau}''}$ with respect to $\tau''(t_o)$:

$$f_{\vec{\tau}'''}(\tau'''_1, \dots, \tau'''_n) = \int_{Su^o(\tau'''_1, \dots, \tau'''_n)} f_{\vec{\tau}''}(\tau''_o, \tau'''_1, \dots, \tau'''_n) d\tau''_o \quad (18)$$

where $Su^o(\tau'''_1, \dots, \tau'''_n)$ is the support of the unknown value $\tau''(t_o)$ when the tuple $\langle \tau''(t_1), \dots, \tau''(t_n) \rangle$ takes the value $\langle \tau'''_1, \dots, \tau'''_n \rangle$. Being a set in DBM form, $D_{t_o}^p$ is convex and thus $Su^o(\tau'''_1, \dots, \tau'''_n)$ is an interval $[Min^o(\tau'''_1, \dots, \tau'''_n), Max^o(\tau'''_1, \dots, \tau'''_n)]$.

By composing Eqs.(18), (16), and (14), we finally express the joint density function of $\vec{\tau}''''$ with respect to that of $\vec{\tau}$:

$$\begin{aligned} f_{\vec{\tau}''''}(\tau'''_1, \dots, \tau'''_n) &= \\ &= \frac{\int_{Min^o(\tau'''_1, \dots, \tau'''_n)}^{Max^o(\tau'''_1, \dots, \tau'''_n)} f_{\vec{\tau}'}(\tau''_o, \tau'''_1 + \tau''_o, \dots, \tau'''_n + \tau''_o) d\tau''_o}{\int_{D_{t_o}^p} f_{\vec{\tau}}(\tau_o, \tau_1, \dots, \tau_n) d\tau_o d\tau_1 \dots d\tau_n} \end{aligned} \quad (19)$$

4. the state probability density function of transitions that are persistent in the child class Σ^c can now be obtained by integrating the density function $f_{\vec{\tau}''''}$ so as to eliminate times to fire of transitions that are not persistent. Specifically, if t_1, \dots, t_m are disabled at the firing, and

$\vec{\tau}_{pers}^c = \langle t_{m+1}, \dots, t_n \rangle$ is the vector of transitions that are persistent in the child class, the density function $f_{\vec{\tau}_{pers}^c}(\tau_{m+1}, \dots, \tau_n)$ is expressed as:

$$f_{\vec{\tau}_{pers}^c}(\tau_{m+1}, \dots, \tau_n) = \int_{-\infty}^{\tau_{m+1}} \dots \int_{-\infty}^{\tau_n} f_{\vec{\tau}''''}(\tau_1, \dots, \tau_m, \tau_{m+1}, \dots, \tau_n) d\tau_1 \dots d\tau_m \quad (20)$$

Finally, the vector $\vec{\tau}''''$ collecting the times to fire of all transitions enabled in the child class is obtained by extending $\vec{\tau}_{pers}^c$ with the vector $\vec{\tau}_{new}^c$ made up by the times to fire of transitions newly enabled in Σ^c , each distributed according to its own static density function.

$$\vec{\tau}'''' = \langle \vec{\tau}_{pers}^c, \vec{\tau}_{new}^c \rangle \quad (21)$$

Since the time to fire of any newly enabled transition t_a is independent from the time fire of any other enabled transition, the joint probability density function $f_{\vec{\tau}''''}$ in the firing domain D^c of the child class Σ^c can be expressed as the product:

$$f_{\langle \vec{\tau}_{new}^c, \vec{\tau}_{pers}^c \rangle}(\vec{\tau}_{new}^c, \vec{\tau}_{pers}^c) = f_{\vec{\tau}_{pers}^c}(\vec{\tau}_{pers}^c) \cdot \prod_{t_a \in \vec{\tau}_{new}^c(\Sigma^c)} f_{t_a}(\tau^c(t_a)) \quad (22)$$

3.4 Example

We illustrate the theory in the derivation of a stochastic class for the example in Fig. 1. For the sake of simplicity, we assume that the firing times of all transitions have a uniform probability density function over their static firing intervals, even if the method can be applied to any kind of distributions.




Figure 1. A stochastic Time Petri Net. All non-deterministic timings are supposed to be uniformly distributed.

Since in the initial class S^0 transitions are newly enabled, their times to fire are all independent. According to this, the joint probability density function over D_0 is obtained as the product of static probability density functions of individual transitions:

$$D_0 = \begin{cases} 5 \leq \tau(t_1) \leq 10 \\ 2 \leq \tau(t_2) \leq 8 \\ 3 \leq \tau(t_3) \leq 9 \end{cases} \quad (23)$$

$$f_0(\tau_1, \tau_2, \tau_3) = \begin{cases} \frac{1}{180} & \text{if } (\tau_1, \tau_2, \tau_3) \in D_0 \\ 0 & \text{if } (\tau_1, \tau_2, \tau_3) \notin D_0 \end{cases}$$

Three events are possible in the class S^0 : the firing of transition t_1 in the interval [5,8], t_2 in [2,8], and t_3 in [3,8]. The assumption of the case that t_3 fires first restricts the firing domain to $D_0^{t_3}$:

$$D_0^{t_3} = \begin{cases} 5 \leq \tau(t_1) \leq 10 \\ 3 \leq \tau(t_2) \leq 8 \\ 3 \leq \tau(t_3) \leq 8 \\ \tau(t_3) \leq \tau(t_1) \\ \tau(t_3) \leq \tau(t_2) \end{cases} \quad (24)$$

According to equation (12), the probability $Prob_{t_3 \text{ first}}$ that t_3 fires first is obtained by integrating $f_0(\tau_1, \tau_2, \tau_3)$ over $D_0^{t_3}$:

$$Prob_{t_3 \text{ first}} = \int_{D_0^{t_3}} f_0(\tau_1, \tau_2, \tau_3) d\tau_1 d\tau_2 d\tau_3 = \frac{29}{90} \quad (25)$$

The joint probability distribution of firing times conditioned to the assumption that t_3 fires is:

$$f_{0|t_3 \text{ first}}(\tau_1, \tau_2, \tau_3) = \begin{cases} \frac{f_0(\tau_1, \tau_2, \tau_3)}{Prob_{t_3 \text{ first}}} = \frac{1}{58} & \text{if } (\tau_1, \tau_2, \tau_3) \in D_0^{t_3} \\ 0 & \text{if } (\tau_1, \tau_2, \tau_3) \notin D_0^{t_3} \end{cases} \quad (26)$$

The class S^1 reached from S^0 through the firing of t_3 has two enabled transitions: t_1 and t_2 . Their firing times are constrained within domain D_1 (also pictured in Fig. 3):

$$D_1 = \begin{cases} 0 \leq \tau(t_1) \leq 7 \\ 0 \leq \tau(t_2) \leq 5 \\ -7 \leq \tau(t_2) - \tau(t_1) \leq 3 \end{cases} \quad (27)$$

According to Eq.(19), we derive the probability density function for transitions t_1 and t_2 by integrating $f_{0|t_3 \text{ first}}(\tau_1'' + \tau_3'', \tau_2'' + \tau_3'', \tau_3'')$ with respect to τ_3'' .

The function $f_{0|t_3 \text{ first}}(\tau_1'' + \tau_3'', \tau_2'' + \tau_3'', \tau_3'')$ is defined over $\hat{D}_0^{t_3}$, derived from $D_0^{t_3}$ through variable substitutions $\tau''(t_3) = \tau(t_3)$, $\tau'''(t_1) = \tau(t_1) - \tau''(t_3)$, $\tau'''(t_2) = \tau(t_2) - \tau''(t_3)$:

$$\hat{D}_0^{t_3} = \begin{cases} 5 \leq \tau'''(t_1) - \tau'''(t_3) \leq 10 \\ 3 \leq \tau'''(t_2) - \tau'''(t_3) \leq 8 \\ 3 \leq \tau'''(t_3) \leq 8 \\ \tau'''(t_1) \geq 0 \\ \tau'''(t_2) \geq 0 \end{cases} \quad (28)$$

In order to integrate $f_{0|t_3 \text{ first}}(\tau_1'' + \tau_3'', \tau_2'' + \tau_3'', \tau_3'')$ with respect to τ_3'' , we must now express the range of variability of $\tau'''(t_3)$ as a function of the values taken by $\tau'''(t_1)$ and $\tau'''(t_2)$. According to Eqs.(18)-(19), this range is an interval $Su^3(\tau_1''', \tau_2''') = [Min^3(\tau_1''', \tau_2'''), Max^3(\tau_1''', \tau_2''')] with:$

$$\begin{aligned} \text{Min}^3(\tau_1''', \tau_2''') &= \min\{\tau_3'' | \langle \tau_1''', \tau_2''', \tau_3'' \rangle \in \widehat{D}_0^{t_3}\} \\ \text{Max}^3(\tau_1''', \tau_2''') &= \max\{\tau_3'' | \langle \tau_1''', \tau_2''', \tau_3'' \rangle \in \widehat{D}_0^{t_3}\} \end{aligned} \quad (29)$$

According to Eq.(28), the two extrema can be expressed as:

$$\begin{aligned} \text{Min}^3(\tau_1''', \tau_2''') &= \max\{5 - \tau_1''', 3 - \tau_2''', 3\} \\ \text{Max}^3(\tau_1''', \tau_2''') &= \min\{10 - \tau_1''', 8 - \tau_2''', 8\} \end{aligned} \quad (30)$$

This splits the range of values for the pair $\langle \tau'''(t_1), \tau'''(t_2) \rangle$ in three sub-regions Z_a, Z_b, Z_c within each of which $\text{Min}^3()$ and $\text{Max}^3()$ has homogeneous form (i.e. it is defined through a single non-piecewise function):

$$\begin{aligned} Z_a &= \begin{cases} 2 \leq \tau'''(t_1) \leq 7 \\ \tau'''(t_2) \geq 0 \\ \tau'''(t_1) - \tau'''(t_2) \geq 2 \end{cases} \\ Z_b &= \begin{cases} 2 \leq \tau'''(t_1) \leq 7 \\ \tau'''(t_2) \leq 5 \\ \tau'''(t_1) - \tau'''(t_2) < 2 \end{cases} \\ Z_c &= \begin{cases} 0 \leq \tau'''(t_1) < 2 \\ \tau'''(t_2) \geq 0 \\ \tau'''(t_1) - \tau'''(t_2) \geq -3 \end{cases} \end{aligned}$$

With reference to this split, we can finally express $\text{Min}^3()$ and $\text{Max}^3()$ as:

$$\begin{aligned} \text{Min}^3(\tau_1''', \tau_2''') &= \begin{cases} 3 & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_a \\ 3 & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_b \\ 5 - \tau_1''' & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_c \end{cases} \\ \text{Max}^3(\tau_1''', \tau_2''') &= \begin{cases} 10 - \tau_1''' & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_a \\ 8 - \tau_2''' & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_b \\ 8 - \tau_2''' & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_c \end{cases} \end{aligned} \quad (31)$$

Fig.2 plots the partitionment for the range of variability of $\langle \tau'''(t_1), \tau'''(t_2) \rangle$ and the form of $\text{Min}^3()$ and $\text{Max}^3()$ in the three subzones. Note that the procedure of derivation of the subzones where $\text{Min}^3()$ and $\text{Max}^3()$ have homogeneous form is general and it is performed as a step in the symbolic computation of the integral in Eq. (19).

The probability density function for transitions t_1 and t_2 is finally derived according to equation (19) and results in a piecewise function defined over the three zones Z_a, Z_b, Z_c (also shown in Fig. 3).

$$f_1(\tau_1''', \tau_2''') = \begin{cases} \frac{1}{58}(7 - \tau_1''') & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_a \\ \frac{1}{58}(5 - \tau_2''') & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_b \\ \frac{1}{58}(3 + \tau_1''' - \tau_2''') & \text{if } \langle \tau_1''', \tau_2''' \rangle \in Z_c \\ 0 & \text{elsewhere} \end{cases} \quad (32)$$

In the child class S^1 reached through t_3 , both t_1 and t_2 are persistent. Moreover, no other transition is newly enabled. According to this, $f_1(\tau_1''', \tau_2''')$ is the probability density function for states collected in the state class S^1 . In the more complex case of any transition being (newly) enabled or disabled by the firing of transition t_3 , we would have to use equations (20) and (22).




Figure 2. The time domain $\widehat{D}_0^{t_3}$ partitioned in three regions during the calculus of $Su^3(\tau_1''', \tau_2''')$; in the projections Z_a, Z_b, Z_c , the bounds of $\tau'''(t_3)$ are both defined by a single homogeneous rule.




Figure 3. The temporal domain D_1 and its state probability density function $f_1(\tau_1, \tau_2)$. D_1 is partitioned in three subzones (a, b and c), representing the three sub-domains of the piecewise function $f_1(\tau_1, \tau_2)$ (see Equation (32)). Note that since the firing of t_3 does not enable or disable any transition, zones a, b and c of D_1 correspond to zones Z_a, Z_b, Z_c reported in Fig. 2.

4 Enumeration

Equations (12) and (22) can be embedded within a "conventional" algorithm for the enumeration of DBM state classes (e.g. [32] [8]) so as to derive a graph of reachability among stochastic state classes of a sTPN.

To this end, algorithms for the detection of class successors and for the computation of their firing domains must be combined with a symbolic derivation of integrals, that can be conveniently supported by a symbolic toolbox. In our experimentation, we integrate the Oris tool for state class enumeration [14] and the Wolfram Mathematica 5.1 for the symbolic calculus [30].

The result of the enumeration of the reachability relation $\Sigma^p \xrightarrow{t, \mu} \Sigma^c$ among stochastic state classes $\Sigma = \langle m, D, f_{\neq}(\cdot) \rangle$ is a stochastic timed transition system, that we call *stochastic class graph*, where nodes are state classes labeled with a

state density function and edges are transitions labeled with a measure of probability.

The stochastic class graph can be regarded as a discrete-time Markov chain X_n with respect to the number n of fired transitions. The analysis of this structure permits to associate a stochastic characterization with symbolic runs identified in the class graph. In particular, it supports the evaluation of such indexes as the probability to reach a logical location, or the probability that the system executes along a given run, the probability that a run exceeds a deadline, the distribution of probabilities for the timing of a symbolic run.




Figure 4. The class graph for the net in Fig.1. Each node is a state class $\langle m, D \rangle$ made up of a marking m and a firing domain D encoded as a difference bounds matrix.

4.1 Example

Enumeration of the reachability relation $S^p \xrightarrow{t} S^c$ among "conventional" state classes $S = \langle m, D \rangle$ for the net in Fig.1 yields the state class graph shown in Fig.4. Markings and time domains for the eight state classes are:

$$\begin{aligned}
 S^0 &= \begin{matrix} 1 & p4 & 1 & p5 & 1 & p6 \\ 5 \leq \tau(t_1) \leq 10 \\ 2 \leq \tau(t_2) \leq 8 \\ 3 \leq \tau(t_3) \leq 9 \end{matrix} & S^1 &= \begin{matrix} 1 & p3 & 1 & p4 & 1 & p5 \\ 0 \leq \tau(t_1) \leq 7 \\ 0 \leq \tau(t_2) \leq 5 \\ -7 \leq \tau(t_2) - \tau(t_1) \leq 3 \end{matrix} \\
 S^2 &= \begin{matrix} 1 & p2 & 1 & p4 & 1 & p6 \\ 0 \leq \tau(t_1) \leq 8 \\ 0 \leq \tau(t_3) \leq 7 \\ -7 \leq \tau(t_3) - \tau(t_1) \leq 4 \end{matrix} & S^3 &= \begin{matrix} 1 & p1 & 1 & p5 & 1 & p6 \\ 0 \leq \tau(t_2) \leq 3 \\ 0 \leq \tau(t_3) \leq 4 \end{matrix} \\
 S^4 &= \begin{matrix} 1 & p1 & 1 & p3 & 1 & p5 \\ 0 \leq \tau(t_2) \leq 3 \end{matrix} & S^5 &= \begin{matrix} 1 & p2 & 1 & p3 & 1 & p4 \\ 0 \leq \tau(t_1) \leq 7 \end{matrix} \\
 S^6 &= \begin{matrix} 1 & p1 & 1 & p2 & 1 & p6 \\ 0 \leq \tau(t_3) \leq 4 \end{matrix} & S^7 &= \begin{matrix} 1 & p1 & 1 & p2 & 1 & p3 \\ 2 \leq \tau(t_4) \leq 4 \end{matrix}
 \end{aligned}$$

When classes are extended with the state density probability, we obtain the extended reachability relation $\Sigma^p \xrightarrow{t, \mu} \Sigma^c$ among stochastic state classes $\Sigma = \langle m, D, f_{\vec{\tau}}(\cdot) \rangle$ shown in Fig. 5. This now includes eleven stochastic state classes as each of the three state classes S^4 , S^5 and S^6 can be reached under two different state probability density functions (see

Fig. 6), thus corresponding to six stochastic state classes ($\Sigma^{4.5}$, $\Sigma^{4.8}$; $\Sigma^{5.7}$, $\Sigma^{5.9}$; $\Sigma^{6.4}$, $\Sigma^{6.6}$).

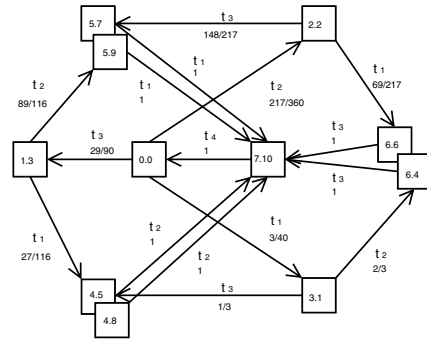


Figure 5. The stochastic class graph for the net in Fig.1. Each node is a stochastic state class $\langle m, D, f_{\vec{\tau}}(\cdot) \rangle$ made up of a marking m , a firing domain D encoded as a difference bounds matrix, and a state density function $f_{\vec{\tau}}(\cdot)$ associating the individual timings within D with a measure of probability. Edges are labeled with a measure of probability associated with the transition. Stochastic state classes are numbered and positioned so as to make evident their correspondence with the classes in the class graph of Fig.4: all stochastic classes labeled by $x.n$ have the marking and the firing domain of class labeled by x , but they differ in the state density function $f_{\vec{\tau}}(\cdot)$.

State density functions for the classes enumerated in the stochastic class graph of Fig.5 are:

$$\begin{aligned}
 f_{0.0} &= \begin{cases} 1/180 & \text{if } 5 \leq \tau_1 \leq 10 \wedge 2 \leq \tau_2 \leq 8 \wedge 3 \leq \tau_3 \leq 9 \\ 0 & \text{elsewhere} \end{cases} \\
 f_{1.3} &= \begin{cases} \frac{1}{58}(7 - \tau_1) & \text{if } 2 \leq \tau_1 \leq 7 \wedge \tau_2 \geq 0 \wedge \tau_1 - \tau_2 \geq 2 \\ \frac{1}{58}(5 - \tau_2) & \text{if } (2 \leq \tau_1 \leq 7 \wedge \tau_2 \leq 5 \wedge \tau_1 - \tau_2 < 2) \\ \frac{1}{58}(3 + \tau_1 - \tau_2) & \text{if } 0 \leq \tau_1 < 2 \wedge \tau_2 \geq 0 \wedge \tau_1 - \tau_2 \geq -3 \\ 0 & \text{elsewhere} \end{cases} \\
 f_{2.2} &= \begin{cases} \frac{10}{217} & \text{if } 2 < \tau_1 \leq 3 \wedge 2 < \tau_1 - \tau_3 \leq 2 \\ -\frac{2}{217}(-8 + \tau_1) & \text{if } \tau_3 > 1 \wedge 3 < \tau_1 \leq 8 \wedge \tau_1 - \tau_3 \geq 1 \\ \frac{2}{217}(3 + \tau_1) & \text{if } 0 \leq \tau_1 \leq 2 \wedge 0 \leq \tau_3 \leq 1 \\ -\frac{2}{217}(-7 + \tau_1 - \tau_3) & \text{if } (\tau_1 - \tau_3 > 2 \wedge \tau_3 \geq 0 \wedge 2 < \tau_1 \leq 3) \\ & \vee (0 \leq \tau_3 \leq 1 \wedge 3 < \tau_1 \leq 7) \\ & \vee (\tau_1 - \tau_3 \leq 7 \wedge \tau_3 \leq 1 \wedge 7 < \tau_1 \leq 8) \\ \frac{2}{217}(4 + \tau_1 - \tau_3) & \text{if } (\tau_1 - \tau_3 \geq -4 \wedge \tau_3 > 1 \wedge 0 \leq \tau_1 < 2) \\ & \vee (-4 \leq \tau_1 - \tau_3 \leq 1 \wedge 2 < \tau_1 \leq 3) \\ -\frac{2}{217}(-7 + \tau_3) & \text{if } 3 < \tau_1 \leq 8 \wedge \tau_1 - \tau_3 < 1 \wedge \tau_3 \leq 7 \\ 0 & \text{elsewhere} \end{cases} \\
 f_{3.1} &= \begin{cases} -\frac{2}{27}(-3 + \tau_2) & \text{if } 0 \leq \tau_2 \leq 3 \wedge \tau_3 \geq 0 \wedge \tau_2 - \tau_3 \geq -1 \\ -\frac{2}{27}(-4 + \tau_3) & \text{if } 0 \leq \tau_2 \leq 3 \wedge \tau_2 - \tau_3 < -1 \wedge \tau_3 \leq 4 \\ 0 & \text{elsewhere} \end{cases}
 \end{aligned}$$

$$f_{6.4} = \begin{cases} \frac{1}{2} & \text{if } 0 \leq \tau_3 < 1 \\ \frac{1}{18}(16 - 8\tau_3 + \tau_3^2) & \text{if } 1 \leq \tau_3 \leq 4 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{4.5} = \begin{cases} \frac{1}{9}(9 - 6\tau_2 + \tau_2^2) & \text{if } 0 \leq \tau_2 \leq 3 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{6.6} = \begin{cases} \frac{1}{23}(13 - 4\tau_3) & \text{if } 0 \leq \tau(\tau_3) < 1 \\ \frac{1}{69}(40 - 14\tau_3 + \tau_3^2) & \text{if } 1 \leq \tau_3 \leq 4 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{5.7} = \begin{cases} \frac{1}{148}(39 + 6\tau_1 - \tau_1^2) & \text{if } 0 \leq \tau_1 \leq 1 \\ \frac{1}{148}(51 - 6\tau_1 - \tau_1^2) & \text{if } 1 < \tau_1 \leq 2 \\ \frac{1}{148}(63 - 16\tau_1 + \tau_1^2) & \text{if } 2 < \tau_1 \leq 7 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{4.8} = \begin{cases} \frac{1}{27}(21 - 10\tau_2 + \tau_2^2) & \text{if } 0 \leq \tau_2 \leq 3 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{5.9} = \begin{cases} \frac{1}{89}(21 + 4\tau_1 - \tau_1^2) & \text{if } 0 \leq \tau_1 < 2 \\ \frac{1}{89}(49 - 14\tau_1 + \tau_1^2) & \text{if } 2 \leq \tau_1 \leq 7 \\ 0 & \text{elsewhere} \end{cases}$$

$$f_{7.10} = \begin{cases} \frac{1}{2} & \text{if } 2 \leq \tau_4 \leq 4 \\ 0 & \text{elsewhere} \end{cases}$$

4.2 Boundedness

Due to the extension of the enumeration algorithm with probabilistic information, the stochastic class graph may include multiple stochastic classes with the same marking and domain but with different state density functions.

The problem is related to confluences occurring at state classes that can be reached through different paths in the class graph, and it can be clearly illustrated with reference to the example net of Fig.1. The class graph in Fig.4 contains a diamond structure made up of four classes S^0 , S^2 , S^3 and S^6 : starting from S^0 , class S^6 can be reached visiting either S^2 (firing transition t_1 and then t_2), or S^3 (vice versa). The ordering of t_1 and t_2 does not influence the set of possible behaviors, but it conditions the distribution of probability in the times to fire of transitions that are enabled in S^6 . In the stochastic class graph of Fig.5, this results in the split of the state class S^6 in two stochastic state classes $\Sigma^{6.6}$ and $\Sigma^{6.4}$, as shown in Fig.6.

The break of confluences in the extension from the class graph to the stochastic class graph not only exacerbates the problem of state space explosion, but may also result in the case of a model which accepts a finite class graph but which has an unbounded stochastic class graph. This condition is related to the existence of cycles in the class graph and to the way in which memory is passed among the transitions that are persistent through the cycle itself.

The case is demonstrated by the infinite overtaking that may occur in the net in Fig.7. The class graph of the net contains a self loop in which transition t_1 fires and re-enables itself leaving t_2 persistent. In the construction of the stochastic reachability graph, the class will be encountered infinite times. In fact, if we assume that t_1 and t_2 have uniform distributions, it can be proved by induction that the state probability density of the stochastic class reached after n subsequent firings of t_1 is equal to:

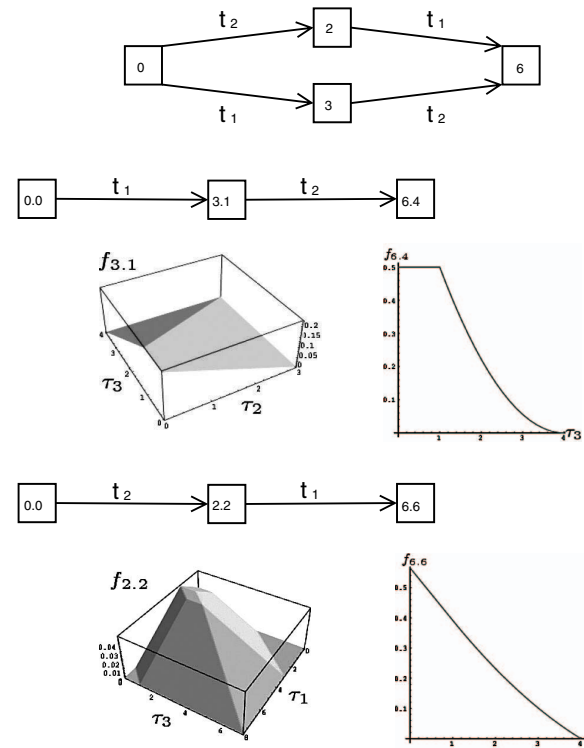


Figure 6. In the class graph of Fig.4, both the timed sequences $S^0 : t_1, t_2$ and $S^0 : t_2, t_1$ lead to the state class S^6 where transition t_3 is constrained to fire in the interval $[0, 4]$. In the stochastic class graph of Fig.4, the same sequences yield two different classes $\Sigma^{6.4}$ and $\Sigma^{6.6}$; the marking and the time domain of these classes are equal, but the state probability density functions are different, reflecting the same range of possibilities with different probabilities.

$$f_{\tau}^n(\tau_1, \tau_2) = \begin{cases} (-1)^n (n+2)(\tau_2 - 1)^{n+1} & \text{if } (\tau_1, \tau_2) \in [0, 1] \times [0, 1] \\ 0 & \text{elsewhere} \end{cases} \quad (33)$$

The example has a clean interpretation: on each firing, transition t_1 re-samples its time to fire within its static firing interval; whereas t_2 always remains persistent and thus accumulates the conditioning of a growing number of events in which it has been overtaken by transition t_1 ; according to this, the density function of t_2 becomes more and more concentrated around the 0 (it tends to the form of a right-Dirac function) and the probability that t_1 overtakes t_2 tends to 0.

It is interesting to note that the accumulation of condi-

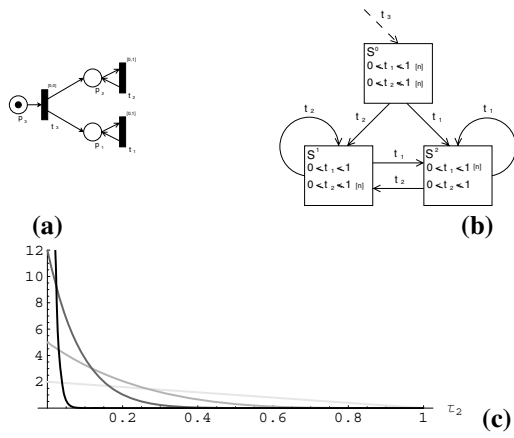


Figure 7. (a) A simple net where the class graph is finite, but the stochastic class graph is unbounded. (b) The class graph of the net includes loops which results in an unbounded number of stochastic state classes, with the same marking and time domain, but with different state probability density functions. (c) In the enumeration of the stochastic class graph, the self loop corresponding to the firing of t_1 from class S^0 yields an unbounded sequence of stochastic classes where the time to fire of transition t_2 is distributed according to a polynomial of increasing order. The picture plots the polynomials generated after 0, 3, 10, and 100 repetitions of the loop.

tioning also occurs along the more elaborate loop in which transitions t_1 and t_2 fire alternatively. In this case, enumeration yields an unbounded sequence of different stochastic classes in which alternatively one of the two transitions is newly enabled (and thus distributed uniformly according to its static density), while the other has a density distributed according to a polynomial of growing order. The form of polynomials is similar to those of Eq.(33) and can be derived through the same kind of procedure reported in the Appendix. In this case, what happens is that the firing transition *passes* its memory to the persistent one through the conditioning that derives from the precedence: starting from the class in which t_1 is newly enabled and t_2 is distributed according to a polynomial of order n , the firing of t_2 yields a new stochastic class in which t_1 becomes distributed according to a polynomial of order $n + 1$.

This observations suggests that unboundedness in the relation between state classes and their associated stochastic state classes is related to the presence of cycles in which each state class has at least one persistent transition that can inherit the conditioning determined by previous firings. According to this, we introduce the following concept:

Definition 4.1 We call *resetting class*, a state class in which all enabled transitions are newly enabled.

By construction, a resetting class is associated with a single stochastic state class in which the times to fire of all enabled transitions are independent and each of them is distributed according to its own static density. This permits to prove the following:

Theorem 4.1 If G is a finite state class graph in which every cyclic path traverses at least one resetting class, then the stochastic class graph Γ associated with G is also finite.

- Ab absurdo, let Γ be unbounded.

Since each stochastic class $\Sigma \in \Gamma$ is associated with a class $S \in G$, there exist a class $S^o \in G$ which is associated with an unbounded number of stochastic classes.

- This implies that the class graph G includes a cyclic path r which originates in S^o , and that the stochastic state graph includes a stochastic class Σ_k^0 associated with S^o , such that, if ρ_k is the path in the stochastic class graph corresponding to r and originating from Σ_k^0 , then the stochastic class Σ_{k+1}^0 reached from Σ_k^0 through the path ρ_k is different than Σ_k^0 :

$$r = S^o \xrightarrow{t_o} S^1 \xrightarrow{t_1} \dots \xrightarrow{t_{N-1}} S^o$$

$$\rho_k = \Sigma_k^0 \xrightarrow{t_o, \mu_k^o} \Sigma_k^1 \xrightarrow{t_1, \mu_k^1} \dots \xrightarrow{t_{N-1}, \mu_k^{N-1}} \Sigma_{k+1}^0 \quad (34)$$

$$\Sigma_k^0 \neq \Sigma_{k+1}^0$$

- Eq.(34) can be easily extended to show that the path ρ_{k+1} which follows the transitions of r starting from the stochastic class Σ_{k+1}^0 visits a sequence of stochastic classes which are all different than the corresponding classes visited along the path ρ_k :

$$\rho_{k+1} = \Sigma_{k+1}^0 \xrightarrow{t_o, \mu_{k+1}^o} \Sigma_{k+1}^1 \xrightarrow{t_1, \mu_{k+1}^1} \dots \xrightarrow{t_{N-1}, \mu_{k+1}^{N-1}} \Sigma_{k+1}^0$$

$$\Sigma_k^n \neq \Sigma_{k+1}^n \quad \forall n = 0, N - 1$$

- Since r is a cyclic path, it visits a resetting class, that we denote as S^{n1} .

Since S^{n1} is visited along r , it is also associated with two stochastic classes Σ_k^{n1} and Σ_{k+1}^{n1} visited along ρ_k and ρ_{k+1} , respectively.

According to Eq.(35), Σ_k^{n1} must be different than Σ_{k+1}^{n1} , which is not possible as Σ_k^{n1} and Σ_{k+1}^{n1} are stochastic classes associated with the same resetting class.

The condition requested for the application of Theorem 4.1 can be easily checked: to this end, it is sufficient considering a reduced class graph G^- which is derived from G by removing every resetting class, and then checking whether G^- includes any cycle. The test can be run in linear time with respect to the size of G and, obviously, without actually constructing the graph G^- . Application of the test gives a positive result (i.e. no unbounded loops are identified) in all the examples of TPN reported in [7] [33] [22]. However, we are actually working on a method which always guarantee boundedness by approximation of state density functions.

5 Conclusions

We have proposed a probabilistic extension of state space analysis for densely timed systems based on time zones encoded through Difference Bounds Matrixes. The approach extends the concept of state class and its reachability relation, commonly applied to the analysis of models such as Time Petri Nets and Timed Automata, by enriching dense firing domains with a state probability density function.

This result, which is the first extension of DBM state classes analysis with probabilistic information, comprises a new approach to bridge the gap between the verification of the possibility of critical behaviors with a quantitative evaluation of their probability.

Acknowledgments

This research was partially supported by the Italian Ministry of Instruction University and Research as a part of the FIRB project *Performance Evaluation of Complex Systems* (PERF).

References

- [1] M. M. Ajmone, G. Balbo, and G. Conte. A class of generalized stochastic petri nets for the performance evaluation of multiprocessor systems. *ACM Trans. on Comp. sys.*, 1984.
- [2] M. M. Ajmone and G. Chiola. On petri nets with deterministic and exponentially distributed firing times. *LNCS*, 266:132–145, 1987.
- [3] R. Alur, C. Courcoubetis, and D. Dill. Model-checking for real-time systems. *Proc. 5th Symp. on Logic in Computer Science*, June 1990.
- [4] R. Alur and D. L. Dill. Automata for modeling real-time systems. *17th ICALP*, 1990.
- [5] R. Alur, T. A. Henzinger, and P. H. Ho. Automatic symbolic verification of embedded systems. *IEEE Trans. on Soft. Eng.*, 22(3), March 1996.
- [6] J. Bengtsson, K. G. Larsen, F. Larsson, P. Pettersson, and W. Yi. *UPPAAL: a Tool-Suite for Automatic Verification of Real-Time Systems*. Hybrid Systems III: LNCS, 1066, 1996.
- [7] B. Berthomieu and M. Diaz. Modeling and verification of time dependent systems using time petri nets. *IEEE Trans. on Soft. Eng.*, 17(3), March 1991.
- [8] B. Berthomieu and M. Menasche. An enumerative approach for analyzing time petri nets. *Proc. IFIP Congress*, 1983.
- [9] A. Bobbio, A. Puliafito, and M. Telek. A modelling framework to implement preemption policies in non-markovian SPNs. *IEEE Trans. On Softw. Eng.*, 26(1), January 2000.
- [10] A. Bobbio and M. Telek. A benchmark for PH estimation algorithms: result for acyclic-PH. *Stochastic models*, 10(661-677), 1994.
- [11] A. Bobbio and M. Telek. Markov regenerative SPN with non-overlapping activity cycles. *Int. Computer Performance and Dependability Symp. - IPDS95*, pages 124–133, 1995.
- [12] J. Bryans. Model checking stochastic automata. *ACM Trans. on Comp. Logic*, pages 452–492, October 2003.
- [13] G. Bucci, L. Sassoli, and E. Vicario. A discrete time model for performance evaluation and correctness verification of real time systems. *Proc. 10th Int. Workshop on Petri Nets and Performance Models*, September 2003.
- [14] G. Bucci, L. Sassoli, and E. Vicario. Oris: a tool for state space analysis of real-time preemptive systems. *Proc. of the 1st Int. Conf. on the Quant. Evaluation of Sys.(QEST)*, 2004.
- [15] H. Choi, V. G. Kulkarni, and K. Trivedi. Transient analysis of deterministic and stochastic petri nets. *Proc. of the 14-th Int. Conf. on Application and theory of Petri Nets*, 1993.
- [16] H. Choi, V. G. Kulkarni, and K. Trivedi. Markov regenerative stochastic petri nets. *Perf. Eval.*, 20:337–357, 1994.
- [17] G. Ciardo, J. Muppala, and K. S. Trivedi. SPNP: stochastic petri net package. *IEEE Int. Workshop on Petri Nets and Performance Models - PNP89*, pages 142–151, 1989.
- [18] C. Daws, A. Olivero, S. Tripakis, and S. Yovine. The tool KRONOS. In *Hybrid Systems III*. 1066, Springer, 1995.
- [19] D. Dill. Timing assumptions and verification of finite-state concurrent systems. *Proc. Workshop on Computer Aided Verification Methods for Finite State Systems*, 1989.
- [20] R. German. *Performance Analysis of Communication Systems with Non-Markovian Stochastic Petri Nets*. John Wiley & Sons, 2000.
- [21] D. Harel and A. Pnueli. On the development of reactive systems. *Logics and models of Concurrent systems*, 1985.
- [22] W.-T. Jong, Y.-S. Shiau, Y.-J. Horng, H.-H. Chen, and S.-M. Chen. Temporal knowledge representation and reasoning techniques using time petri nets. *IEEE Trans. on Soft. Eng., Man and Cybernetics, NATO*, 29(4), August 1999.
- [23] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Verifying quantitative properties of continuous probabilistic timed automata. *Proc. CONCUR'00*, 1877:123–137, 2000.
- [24] D. Lime and O. H. Roux. Expressiveness and analysis of scheduling extended time petri nets. In *5th IFAC conference on fieldbus and their applications (FET 2003)*, Portugal, 2003. Elsevier Science.
- [25] G. I. I. Lopez, H. Hermanns, and J.-P. Katoen. Beyond memoryless distributions: Model checking semi-markov chains. In *Proc. PAFM-ProbMiV 2001*, 2165:57–70, LNCS 2001.
- [26] M. K. Molloy. Discrete time stochastic petri nets. *IEEE Trans. on Soft. Eng.*, 11(1), January 1985.
- [27] M. F. Neuts. *Matrix Geometric Solutions in Stochastic Models*. Johns Hopkins University Press, 1981.
- [28] W. Penczek and A. Polrola. Specification and model checking of temporal properties in time petri nets and timed automata. *Proceedings of the 25th Int. Conf on Application and Theory of Petri Nets, ICATPN2004*, June 2004.
- [29] K. G. Popstojanova and K. S. Trivedi. Stochastic modeling formalisms for dependability, performance and performability. *Performance Evaluation - Origins and Directions, LNCS*, pages 385–404, 2000.
- [30] W. Research. Mathematica 5.1. www.wolfram.com.
- [31] J. A. Stankovic and K. Ramamritham. What is predictability for real time systems. *Journal of Real Time Systems*, 2, 1990.
- [32] E. Vicario. Static analysis and dynamic steering of time dependent systems using time petri nets. *IEEE Trans. on Soft. Eng.*, August 2001.
- [33] D. Xu, X. He, and Y. Deng. Compositional schedulability analysis of real-time systems using time petri nets. *IEEE Trans. on Soft. Eng.*, 28(10), October 2002.