©2009 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.

# Using Stochastic State Classes in Quantitative Evaluation of Dense-Time Reactive Systems

Enrico Vicario, *Member*, *IEEE Computer Society*, Luigi Sassoli, and Laura Carnevali

**Abstract**—In the verification of reactive systems with nondeterministic densely valued temporal parameters, the state-space can be covered through equivalence classes, each composed 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 events ordering and stimulus/response deadlines, but it does not provide any measure of probability for feasible behaviors. We extend DBM equivalence classes with a density-function which provides a measure for the probability of individual states. To this end, we extend Time Petri Nets by associating a probability density-function to the static firing interval of each nondeterministic transition. We then explain how this stochastic information induces a probability distribution for the states contained within a DBM class and how this probability evolves in the enumeration of the reachability relation among classes. This enables the construction of a stochastic transition system which supports correctness verification based on the theory of TPNs, provides a measure of probability for each feasible run, enables steady-state analysis based on Markov Renewal Theory. In so doing, we provide a means to identify feasible behaviors and to associate them with a measure of probability in models with multiple concurrent generally distributed nondeterministic timers.

**Index Terms**—Correctness verification, performance and dependability evaluation, stochastic Time Petri nets, non-Markovian Stochastic Petri Nets, dense-time state-space analysis, Difference Bounds Matrix, Markov Renewal Theory.

## 1 Introduction

In the construction of concurrent timed systems, the engineering of reliability and efficiency characteristics [27] often intertwines verification of sequencing and timeliness correctness [26], [39] with the evaluation of quantitative measures of performance and dependability [36]. But, despite this need of the application context, qualitative identification of feasible behaviors and quantitative evaluation of their probability usually rely on separate formalisms and techniques. In particular, integration of the two aspects for densely timed models is an open and widely addressed problem.

In a large part of the literature on quantitative performance and dependability evaluation, stochastic temporal parameters are associated with a negative exponential distribution. Notable examples are Stochastic Petri Nets (SPN) and Generalized Stochastic Petri Nets (GSPN) [1], [33], which enable Markovian analysis [17] but prevent representation of implicit precedences induced by finite timing constraints (e.g., time-outs or firm schedules). Various techniques for the analysis of non-Markovian SPN break the limit of exponential timing [23], [10], but always with relevant restrictions in the representation of behaviors where concurrency is determined by timing constraints. Deterministic and Stochastic Petri Nets (DSPN) [2] and Markov Regenerative Stochastic Petri Nets (MRSPN) [15], [12], [24], [16] include transitions with deterministic or generally distributed bounded delay, respectively. If the model is constructed so as to avoid having multiple generally distributed transitions concurrently enabled, the set of reachable markings is isomorphic to that of the untimed model and the underlying process can be analyzed through Markov Renewal Theory [16], [15] or through the method of supplementary variables [24], [25], [40]. The limitation on concurrent enabling is overcome in [31], [32] through the analysis of a discrete time General-State-Space Chain embedded in the stochastic process at equidistant time points, but the analysis relies on the assumption that all generally distributed transitions are deterministic. In [11], transition delays are represented through continuous Phase-Type distributions [34] and analysis is carried out by combining each reachable marking of the untimed model with the phases of the distributions of enabled transitions. The use of Kronecker algebra circumvents explicit expansion of the timed state-space [34] provided that the reachability relation of the untimed marking is bounded. However, the formalism cannot represent firmly bounded delays, thus preventing again the representation of sequencing constraints induced by mutual delays among concurrent computations or clocks.

In the context of correctness verification, a number of analysis techniques have been proposed for formalisms such as Time Petri Nets (TPN) and Timed Automata (TA), which include nondeterministic temporal parameters taking values within (finite or unbounded) dense intervals [7], [41], [4], [6], [18]. For this kind of model, the timed state-space is covered through the enumeration of a discrete reachability relation among state-classes, each composed of a discrete logical location and a time-domain containing a dense variety of timings. In particular, a wide literature has developed upon state-classes that encode time-domains as the solution space of a set of linear inequalities constraining

• *The authors are with the Dipartimento Sistemi e Informatica, Università di Firenze, 3, via santa Marta, 50139 Firenze, Italy.*
*E-mail: {vicario, sassoli, carnevali}@dsi.unifi.it.*

*Manuscript received 22 Dec. 2006; revised 7 Dec. 2007; accepted 21 Jan. 2009; published online 20 May 2009.*
*Recommended for acceptance by S. Donatelli.*
*For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number TSE-0291-1206.*
*Digital Object Identifier no. 10.1109/TSE.2009.36.*

098-5589/09/$25.00 © 2009 IEEE Published by the IEEE Computer Society

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

the difference between pairs of clocks, also known as Difference Bounds Matrixes (DBM) [3], [8], [7], [19], [5]. DBM domains enable efficient representation and manipulation and encompass in the same class the dense set of states that can be reached through runs that differ in the timing but share a common qualitative ordering of events. Enumeration of the reachability relation among DBM state-classes thus provides a means to solve problems of reachability and model checking under the restrictions induced by timing constraints [7], [6], [30], [41] and to evaluate min-max constraints on the time spent along any symbolic run [41]. However, analysis techniques based on DBM state-classes do not associate reachable states and feasible behaviors with any measure of probability, thus preventing application in performance evaluation and in those contexts where qualitative correctness verification is merged with quantitative dependability evaluation.

In this paper, we address the integration of qualitative verification and quantitative evaluation based on a probabilistic extension of DBM theory. To this end, we extend Time Petri Nets into stochastic Time Petri Nets (sTPN) by decorating the static firing interval of each transition with a generally distributed probability density-function and with a competitiveness weight (Section 2). We then explain how this stochastic information induces a measure of probability for individual states contained in a DBM state-class and how this probability evolves in the enumeration of the reachability relation among state-classes (Sections 3 and 4). This enables enumeration of a stochastic transition system, called stochastic-graph, which extends verification of qualitative correctness properties with the evaluation of the probability of any selected symbolic run and of the distribution of probability of its duration (Section 5). We then provide a characterization of the stochastic process underlying the evolution of an sTPN model and present a technique for steady-state analysis based on Markov Renewal Theory (Section 5.3). We finally illustrate through a case example the kind of results that can be derived through the proposed technique, its present limits of applicability, and its relation with previous techniques (Sections 6 and 7). For the sake of readability, all Proofs are reported in Appendix (Section 8).

## 2 Stochastic Time Petri Nets

A stochastic Time Petri Net (sTPN) is a tuple

$$sTPN = \langle P; T; A^-; A^+; m_0; EFT^s; LFT^s; \tau_0; \mathcal{F}; \mathcal{C} \rangle. \quad (1)$$

$P$ is a set of places, $T$ a set of transitions disjoint from $P$, $A^- \subseteq P \times T$ and $A^+ \subseteq T \times P$ are sets of precondition and postcondition arcs. Place $p$ is said to be an input or an output place for transition $t$ if $\langle p, t \rangle \in A^-$ or $\langle t, p \rangle \in A^+$, respectively. $m_0 : P \to \mathbb{N}$ is the (initial) marking associating each place with a nonnegative number of tokens. $EFT^s$ and $LFT^s$ associate each transition $t \in T$ with an earliest and a (possibly infinite) latest static firing time:

$$EFT^s : T \to \mathbb{R}_0^+ \quad LFT^s : T \to \mathbb{R}_0^+ \cup \{+\infty\}, \quad (2)$$

where $\mathbb{R}_0^+$ denotes the set of nonnegative real numbers and $EFT^s(t) \leq LFT^s(t)$. Besides, $\tau_0$ associates each transition $t \in T$ with an (initial) time-to-fire: $\tau_0 : T \to \mathbb{R}_0^+$. $\mathcal{F}$ and $\mathcal{C}$ define a measure of probability for nondeterministic choices: $\mathcal{C}$ associates each transition with a weight $\mathcal{C}$: $T \to \mathbb{R}^+$, where $\mathbb{R}^+$ denotes the set of positive (nonnull) real numbers; $\mathcal{F}$ associates each transition $t \in T$ with a static probability distribution $F_t^s()$ supported in the static firing interval $[EFT^s(t), LFT^s(t)]$ (i.e., $x < EFT^s(t) \to F_t^s(x) = 0$, $x > EFT^s(t) \to F_t^s(x) > 0$, $x \geq LFT^s(t) \to F_t^s(x) = 1$). If $EFT^s(t) \neq LFT^s(t)$, we assume that $F_t^s()$ is absolutely continuous and, thus, that there exists a density-function $f_t^s()$ such that:

$$F_t^s(x) = \int_0^x f_t^s(y)dy. \quad (3)$$

The *state* of an sTPN is a pair $s = \langle m, \tau \rangle$, where $m : P \to \mathbb{N}$ is a marking and $\tau : T \to \mathbb{R}_0^+$ associates each transition with a (dynamic) time-to-fire. The state evolves according to two clauses of firability and firing, which extend the clock semantics of Time Petri Nets with a stochastic characterization of nondeterminism.

*Firability:* A transition $t_o$ is *enabled* if each of its input places contains at least one token and it is *firable* if it is enabled and its time-to-fire $\tau(t_o)$ 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 $\mathcal{C}$:

$$Prob\{t_0 \text{ is selected}\} = \frac{\mathcal{C}(t_o)}{\sum_{t_i \text{ firable}} \mathcal{C}(t_i)}. \quad (4)$$

*Firing:* When $t_o$ fires, the state $s = \langle m, \tau \rangle$ is replaced by a new state $s' = \langle m', \tau' \rangle$, which we write as $s \xrightarrow{t_o} s'$. 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 $t_o$:

$$\begin{array}{ll} m_{tmp}(p_{in}) = m(p_{in}) - 1, & \forall p_{in}.\langle p_{in}, t_o \rangle \in A^-, \\ m'(p_{out}) = m_{tmp}(p_{out}) + 1, & \forall p_{out}.\langle t_o, p_{out} \rangle \in A^+. \end{array} \quad (5)$$

Transitions that are enabled by the intermediate marking $m_{tmp}$ (and thus by the final marking $m'$) are *persistent*, while those that are enabled by $m'$ but not by $m_{tmp}$ are *newly enabled*. As in [7], [41], if $t_o$ is still enabled after its own firing, it is regarded as newly enabled. For any transition $t_i$ that is persistent after the firing of $t_o$, the time-to-fire is reduced by the time elapsed in the previous state:

$$\tau'(t_i) = \tau(t_i) - \tau(t_o), \quad (6)$$

whereas, for any transition $t_a$ that is newly enabled after the firing of $t_o$, the time-to-fire takes a random value sampled according to the static probability distribution $F_{t_a}^s()$, thus falling within the static firing interval of $t_a$:

$$\begin{array}{l} EFT^s(t_a) \leq \tau'(t_a) \leq LFT^s(t_a), \\ Prob\{\tau'(t_a) \leq x\} = F_{t_a}^s(x). \end{array} \quad (7)$$

## 3 From States to Stochastic Classes

### 3.1 States and Classes

A newly enabled transition may take any real-valued time-to-fire within its static firing interval, and each value leads to

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

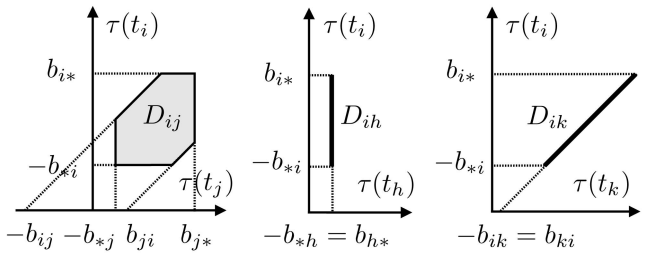


Fig. 1. Three bivariate DBM domains with different dimensions. Since $\|D_{ij}\| > 0$, both $\tau(t_i)$ and $\tau(t_j)$ are distributed. $\|D_{ih}\| = 0$ as $\tau(t_h)$ is bound to assume a deterministic value. $\|D_{ik}\| = 0$ as $\tau(t_k) - \tau(t_i)$ is bound to assume a deterministic value.

a different state, yielding a dense variety of possible successors. To obtain a discrete representation of the state-space, the reachability relation between states is replaced by a symbolic reachability relation between *state-classes*. Each state-class, called class in the following, is a dense set of states with the same marking but with different values for times to fire of enabled transitions. This is represented as a pair:

$$Class = \langle m, D \rangle \quad D \subseteq \mathbb{R}^N, \tag{8}$$

where $m$ is a marking, $D$ is a *time-domain* $D$, $N$ is the cardinality of the set $|T^e(m)|$ of transitions enabled by $m$. The structure of $D$ jointly depends on the firing clause that defines how timers advance and on the semantics of the reachability relation among classes. Most works (and this among them) on the analysis of densely timed models are based on the AE reachability relation [35], [41].

**Definition 3.1.** *Class $S'$ is a successor of class $S$ through $t_o$, which we write as $S \xrightarrow{t_o} S'$, if and only if $S'$ contains all and only the states that are reachable from some state contained in $S$ through some feasible firing of $t_o$.*

Under this reachability relation, a time-domain $D$ can be represented as the set of solutions of a set of linear inequalities in the form of a DBM [19]:

$$D = \begin{cases} \tau(t_i) - \tau(t_j) \leq b_{ij} \\ \forall t_i, t_j \in T^e(m) \cup \{t_*\} \quad \text{with } t_i \neq t_j, \end{cases} \tag{9}$$

where $\tau(t_*) = 0$ is a fictitious variable denoting the ground time of the class and $b_{ij} \in \mathbb{R} \cup \{+\infty\}$ are the coefficients that define the boundaries of the time-domain (see Fig. 1).

The DBM form has a normal representation where $b_{ij}$ coincides with the maximum feasible value that can be attained by the difference $\tau(t_i) - \tau(t_j)$. If the set of solutions is not empty, the normal form exists uniquely and is univocally identified by the condition:

$$b_{ij} \leq b_{ih} + b_{hj} \quad \begin{array}{l} \forall t_i, t_j, t_h \in T^e(m) \cup \{t_*\} \\ \text{with } t_i \neq t_j \neq t_h \neq t_i. \end{array} \tag{10}$$

According to (10), the normal form is the solution of an all-shortest-path problem and can be computed in time $O(N^3)$ or even in time $O(N^2)$ in a repeated derivation exploiting warm restart [41]. The normal form supports efficient detection and derivation of successor classes, in time $O(N)$ and $O(N^2)$, respectively [41].

### 3.2 Stochastic Classes

The class succession relation $S \xrightarrow{t_0} S'$ enables the analysis of the range of feasible behaviors of a model [41]. In order to associate these behaviors with a quantitative measure of probability, we extend classes into *stochastic-classes* through the addition of a probability distribution over the time-domain:

$$Stochastic\ class = \langle m, D, F_\tau() \rangle, \tag{11}$$

where $m$ is a marking, $D$ is a DBM domain for times to fire of transitions enabled by $m$, and $F_\tau()$ is a probability distribution supported in $D$. Since each determination of $\tau$ uniquely identifies a state in $S$ and vice versa, $F_\tau()$ takes the meaning of a probability distribution for the states contained in $S$, and we thus call it *state probability distribution*. With this perspective, we introduce a stochastic reachability relation.

**Definition 3.2.** *Given two stochastic-classes $\Sigma = \langle m, D, F_\tau() \rangle$ and $\Sigma' = \langle m', D', F_{\tau'}() \rangle$, we say that $\Sigma'$ is a successor of $\Sigma$ through $t_o$ with probability $\mu$, and we write $\Sigma \xrightarrow{t_o, \mu} \Sigma'$, 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 $\tau$ with support $D$ and distribution $F_\tau()$, then the firing of $t_o$ can occur with probability $\mu$ and leads to a new marking $m'$ and a new vector of times to fire $\tau'$ with support $D'$ and distribution $F_{\tau'}()$.*

### 3.3 DBM Partitioned Form

The time-domain $D$ of (9) may have a null measure in $\mathbb{R}^N$. This occurs if and only if: 1) There exists an enabled transition $t_h$ such that $b_{h*} + b_{*h} = 0$, which means that $t_h$ is bound to fire after a deterministic delay $b_{h*}$ since the time of entrance in the class (see $D_{ih}$ in Fig. 1), or 2) there exist two enabled transitions $t_i$ and $t_k$ such that $b_{ik} + b_{ki} = 0$, which means that $t_k$ is bound to have a deterministic delay $b_{ik}$ with respect to $t_i$ (see $D_{ik}$ in Fig. 1). In both cases, the stochastic characterization of the space of solutions of the time-domain can be limited to a minimal set of components that are sufficient to determine all the unknown values of $D$. This reduces the number of managed variables limiting spatial complexity in the representation of DBM domains. More importantly for the present treatment, this avoids a stochastic characterization for deterministic or deterministically dependent variables. With this aim, we partition the set of times to fire $\tau(t_n)$ of (9) into three subsets $Q$, $U$, and $V$:

- $Q$ is a set of maximal cardinality that includes only times to fire that are nondeterministic and have a nondeterministic mutual delay (we call them *distributed*):

$$\begin{array}{ll} \tau(t_i) \in Q \text{ iff} & (b_{i*} + b_{*i} \neq 0) \wedge \\ & (\forall \tau(t_j) \in Q \text{ with } i \neq j, b_{ij} + b_{ji} \neq 0). \end{array} \tag{12}$$

- $U$ includes the times to fire $\tau(t_h)$ that are bound to take a deterministic value (we call them *deterministic*):

$$\tau(t_h) \in U \quad \text{iff} \quad b_{h*} + b_{*h} = 0. \tag{13}$$

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

- $V$ includes the times to fire that are bound to have a deterministic delay with respect to some time-to-fire in $Q$ (we call them *synchronized*):

$$\tau(t_k) \in V \text{ iff } \begin{array}{l}(b_{k*} + b_{*k} \neq 0) \wedge \\ (\exists \tau(t_{\iota_k}) \in Q.\ b_{k\iota_k} + b_{\iota_k k} = 0).\end{array} \quad (14)$$

For example, with reference to Fig. 1: In $D_{ij}$, $\langle Q, U, V\rangle = \langle\{\tau(t_i), \tau(t_j)\}, \phi, \phi\rangle$; in $D_{ih}$, $\langle Q, U, V\rangle = \langle\{\tau(t_i)\}, \{\tau(t_h)\}, \phi\rangle$; in $D_{ik}$, $\langle Q, U, V\rangle = \langle\{\tau(t_i)\}, \phi, \{\tau(t_k)\}\rangle$. Note that, while $U$ is always uniquely determined, the partition between $Q$ and $V$ is not unique: For every $\tau(t_k) \in V$, the elements $\tau(t_k)$ and $\tau(t_{\iota_k})$ can be swapped between $V$ and $Q$. For instance, the set of variables of $D_{ik}$ of Fig. 1 could also be partitioned as $\langle Q, U, V\rangle = \langle\{\tau(t_k)\}, \phi, \{\tau(t_i)\}\rangle$.

By referring to the $\langle Q, U, V\rangle$ partition, the normal form of the set $D$ of (9) can be rewritten in a form, which we call *partitioned*, that combines the set $D_Q$ of constraints limiting the elements of $Q$ with a set of equalities expressing all the remaining variables of $U$ and $V$ as a function of the solution of $D_Q$:

$$D_Q = \begin{cases} \tau(t_i) - \tau(t_j) \leq b_{ij} \\ \tau(t_i) - \tau(t_*) \leq b_{i*} \\ \tau(t_*) - \tau(t_i) \leq b_{*i} \\ \forall \tau(t_i), \tau(t_j) \in Q. \end{cases}$$
$$D = \begin{cases} D_Q \\ \tau(t_h) = b_{h*} + \tau(t_*) \\ \tau(t_k) = b_{k\iota_k} + \tau(t_{\iota_k}) \\ \forall \tau(t_h) \in U, \quad \forall \tau(t_k) \in V. \end{cases} \quad (15)$$

By construction, $D_Q$ has a nonnull measure in $\mathbb{R}^{|Q|}$.

### 3.4 Stochastic Classes in Partitioned Form

When the DBM domain $D$ is represented in partitioned form, the state distribution $F_\tau()$ is completely defined by the state-density-function $f_{\tau_Q}()$ associated with the variables in $Q$:

$$F_\tau(X) = \int_{D(X)\downarrow Q} f_{\tau_Q}(x)dx, \quad (16)$$

where $X$ is a point in $Q \times U \times V$, $D(X)$ is the subset of $D$ upper-bounded by $X$ (i.e., $D(X) = \{\chi \in Q \times U \times V | \chi_n \leq X_n \forall n \in [0,\ N-1]\}$ with $N = |Q \times U \times V|$), and $\Delta \downarrow Q$ denotes the projection of a set $\Delta \subseteq Q \times U \times V$ on the subspace $Q$:

$$\Delta \downarrow Q \stackrel{def}{=} \{q \in Q |\ \exists u \in U, \exists v \in V \text{ such that } \langle q, u, v\rangle \in \Delta\}. \quad (17)$$

Since $f_Q()$ is sufficient to determine $F_\tau()$, we will represent a stochastic-class $\langle m, D, F_\tau()\rangle$ as $\langle m, D, f_Q()\rangle$.

## 4 Enumeration

### 4.1 Successors Detection

The firing of transition $t_o$ is a *possible outcoming event* from a stochastic-class $\Sigma = \langle m, D, F_\tau()\rangle$ iff $t_o$ is enabled by $m$ and $D$ accepts solutions such that $t_o$ has a time-to-fire not higher than that of any other enabled transition (i.e., $\forall t_n \in T^e(m)$, $\tau(t_o) \leq \tau(t_n)$). This occurs iff there is a nonempty set of solutions for the following *restricted firing domain* $D_{t_o}$:

$$D_{t_o} = \begin{cases} D \\ \tau(t_o) - \tau(t_i) \leq 0 \\ \tau(t_o) \leq b_{h*} \\ \tau(t_o) - \tau(t_{\iota_k}) \leq -b_{\iota_k k} \\ \forall \tau(t_i) \neq \tau(t_j) \in Q,\ \forall \tau(t_h) \in U,\ \forall \tau(t_k) \in V, \end{cases} \quad (18)$$

where $\tau(t_{\iota_k})$ is the distributed time-to-fire synchronized with $\tau(t_k)$ (as in (14)).

According to [41]: $D_{t_o}$ accepts solutions iff $b_{no} \geq 0$ for every enabled transition $t_n \neq t_o$; in this case, the normal form of $D_{t_o}$, whose coefficients will be denoted by $B_{ij}$, can be derived from the normal form of $D$ in time $O(N^2)$.

### 4.2 Calculus of Successors Probability

Let $t_o$ be a possible outcoming event and denote its probability by $\mu_{t_o}$. If all enabled transitions are deterministic (i.e., if $Q = \phi$), then $D$ accepts a single solution $\tau$. Since $t_o$ is assumed to be a possible outcoming event, its time-to-fire is minimal. Any other enabled transition $t_n$ is thus a possible outcoming event iff $\tau_n = \tau_o$. Since both $t_n$ and $t_o$ are deterministic and $D$ is represented in normal form, this occurs iff $b_{no} = b_{on} = 0$. According to this, $\mu_{t_o}$ is the probability that $t_o$ is selected in the random switch among $t_o$ and all and only the transitions such that $b_{no} = b_{on} = 0$:

$$\mu_{t_o} = \frac{\mathcal{C}(t_o)}{\mathcal{C}(t_o) + \sum_{\substack{t_n \neq t_o \in T^e(m) \\ with\ b_{no} = b_{on} = 0}} \mathcal{C}(t_n)}. \quad (19)$$

If also distributed transitions are enabled (i.e., if $Q \neq \phi$), then $\mu_o$ is the joint probability that the vector of times to fire $\tau$ takes a value such that $\tau_o \leq \tau_n$ for every enabled transition $t_n$ (which is true iff $\tau \in D_{t_o}$) and $t_o$ is selected in the random switch among all of the transitions $t_n \in T^e(m)$ such that $\tau_n = \tau_o$. According to this, $\mu_{t_o}$ can be expressed as

$$\mu_{t_o} = \int_{D_{t_o}\downarrow Q} \frac{\mathcal{C}(t_o)}{\mathcal{C}(t_o) + \sum_{t_n \in T^f(\Omega)\setminus\{t_o\}} \mathcal{C}(t_n)} \cdot f_{\tau_Q}(\Omega)d\Omega, \quad (20)$$

where $T^f(\Omega)$ is the set of transitions whose time-to-fire is minimal when the projection of the vector of times to fire over $Q$ is $\Omega$. If $b_{no} = b_{on} = 0$, then $\tau_n = \tau_o$ all over $D$; in this case, the time-to-fire of $t_n$ is minimal all over $D_{t_o}$, and $t_n \in T^f(\Omega)$ for every $\Omega \in D_{t_o} \downarrow Q$. On the other hand, if $b_{no} \neq 0$ or $b_{no} \neq b_{on}$, then the projection over $Q$ of the restriction of $D$ with the additional constraint $\tau_o = \tau_n$ has a null measure. According to this, almost everywhere in $D_{t_o} \downarrow Q$, the set $T^f(\Omega)$ coincides with the set of enabled transitions $t_n$ such that $b_{no} = b_{on} = 0$. Since this condition does not depend on $\Omega$, the random switch of (20) can be taken out of the integral, yielding:

$$\mu_{t_o} = \frac{\mathcal{C}(t_o)}{\mathcal{C}(t_o) + \sum_{\substack{t_n \neq t_o \in T^e(m) \\ with\ b_{no} = b_{on} = 0}} \mathcal{C}(t_n)} \cdot \int_{D_{t_o}\downarrow Q} f_{\tau_Q}(\Omega)d\Omega. \quad (21)$$

It is relevant to note that it may be the case that $t_o$ is possible but has null probability (i.e., $D_{t_o} \neq \phi \wedge \mu_{t_o} = 0$). According to (19) and (21), this occurs iff the set $Q$ of distributed transitions is not empty and the projection of the restricted domain $D_o$ over $Q$ has a null measure (see also Fig. 2):

$$(D_{t_o} \neq \phi \wedge \mu_{t_o} = 0) \text{ iff } (Q \neq \phi \wedge \|D_{t_o} \downarrow Q\| = 0). \quad (22)$$

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

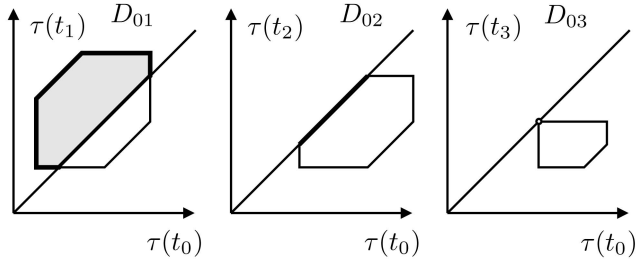


Fig. 2. Three bivariate DBM domains illustrating how the restriction $\tau(t_o) \leq \tau(t_i)$ may change the $\langle Q, U, V\rangle$ partition. For $D_{01}$, the restricted domain (drawn in gray) maintains a nonnull measure, both $\tau(t_0)$ and $\tau(t_1)$ remain distributed, and $\mu_{t_0} > 0$. For $D_{02}$, $\tau(t_0)$ and $\tau(t_2)$ become synchronized to each other. For $D_{03}$, both $\tau(t_0)$ and $\tau(t_3)$ become deterministic. In both cases of $D_{02}$ and $D_{03}$, the firing of $t_o$ is possible but has null probability.

In the enumeration of the stochastic relation $\overset{t_o,\mu_{t_o}}{\Longrightarrow}$, we suppress the edges with $\mu_{t_o} = 0$ (i.e., we consider only events with non-null probability $\mu_{t_o} > 0$). These edges are usually considered in the practice of qualitative analysis [41], [7] where events occurring over a time-domain with null measure may serve to detect critical behaviors enabled by boundary conditions, and where in any case there is not a measure of probability (neither assumptions on its regularity) that permits to tell that a null measure domain is reached with null probability. In a technical perspective, the present treatment could encompass also firings with null probability, but these would not affect reward measures and would yield badly conditioned numerical problems in Markovian analysis.

### 4.3 Successor Derivation

We denote by $\Sigma' = \langle m', D', f_{Q'}()\rangle$ the stochastic-class reached from $\Sigma = \langle m, D, f_Q()\rangle$ through $t_o$ (i.e., $\Sigma \overset{t_o,\mu_{t_o}}{\Longrightarrow} \Sigma'$). Marking $m'$ is derived by moving tokens according to (5). Besides, $D'$ and $f_{Q'}()$ are derived through four steps, accounting for

1. the conditioning introduced by the assumption that $t_o$ is the first transition to fire;
2. the shift of times to fire observed at the firing of $t_o$;
3. the projection that eliminates from the set of unknown values the time-to-fire of any transition disabled by the firing;
4. the addition of times to fire of newly enabled transitions.

The four steps are developed under the assumption that $\tau(t_o)$ is either distributed or deterministic; the case in which $\tau(t_o)$ is synchronized can be encompassed in the treatment by swapping $\tau(t_o)$ and $\tau(t_{\iota_o})$ in the partition of the sets $Q$ and $V$.

#### 4.3.1 Conditioning

The assumption that $t_o$ is the next transition to fire conditions the vector $\tau$, yielding a new random variable $\tau^a$:

$$\tau^a = \tau | \tau_o \leq \tau_n \ \forall \tau_n \in T^e(m). \tag{23}$$

The assumption $\mu_{t_o} \neq 0$ guarantees that the conditioning does not change the $\langle Q, U, V\rangle$ partition.

**Lemma 4.1.** *If $\mu_{t_o} > 0$, then the partition $\langle Q, U, V\rangle$ is not changed by the conditioning induced by the assumption that $\tau_o$ is minimum.*

By construction, $\tau^a$ is supported over the restricted domain $D_{t_o}$ and its density-function $f_{Q^a}()$ can be derived as a conditional probability:

$$f_{Q^a}(x) = \frac{f_Q(x)}{\mu_{t_o}}. \tag{24}$$

#### 4.3.2 Time Advancement

At the firing of $t_o$, time is observed to advance by the value of $\tau^a(t_o)$. This yields a new vector of times to fire $\tau^b$ where the times to fire of all transitions other than $t_o$ are reduced by the value of $\tau^a(t_o)$, and $\tau^a(t_o)$ itself is eliminated:

$$\tau^b(t_n) = \tau^a(t_n) - \tau^a(t_o) \quad \forall t_n \in T^e(m) \backslash \{t_o\}. \tag{25}$$

The effect of the advancement is basically different whether $t_o$ is deterministic or distributed. If $\tau^a(t_o)$ is deterministic, then $\tau^a(t_o) = B_{o*} = B_{*0}$. In this case, the $\langle Q, U, V\rangle$ partition is not changed, the support $D_{Q^b}$ of distributed times to fire is obtained through a constant shift of $D_{Q^a}$, and the joint density-function $f_{Q^b}()$ within $D_{Q^b}$ is expressed as

$$f_{Q^b}(x^b) = f_{Q^a}(x_1 + B_{o*}, \ldots, x_{|Q^a|-1} + B_{o*}). \tag{26}$$

Whereas, if $\tau(t_o)$ is distributed, the $\langle Q, U, V\rangle$ partition changes: If $\tau^a(t_i) \in Q^a$, then $\tau^b(t_i) \in Q^b$ (i.e., distributed times to fire remain distributed); if $\tau^a(t_n) \in V^a$ with $\iota_n = 0$ (i.e., if $\tau^a(t_n)$ was synchronized with $\tau^a(t_o)$), then $\tau^b(t_n) = B_{no}$ and thus $\tau^b(t_n)$ becomes deterministic; if $\tau^a(t_k) \in V^a$ with $\iota_k \neq 0$ (i.e., if $\tau^a(t_k)$ was synchronized with a time-to-fire other than $\tau^a(t_o)$), then $\tau^a(t_k)$ remains synchronized with $\tau^b(t_{\iota_k})$ with the same initial delay $B_{kt_k}$; all times to fire that were deterministic become distributed, but they remain synchronized among themselves, and, thus, one of them is added to $Q^b$ and the other ones are added to $V^b$.

The density-function $f_{Q^b}()$ takes a different form whether $U^a = \phi$ or $U^a \neq \phi$ (i.e., whether the initial partition $\langle Q^a, U^a, V^a\rangle$ contains any deterministic transition). If $U^a = \phi$, then $Q^b$ is derived from $Q^a$ by removing $\tau^a(t_o)$, the support $D_{Q^b}$ of $f_{Q^b}()$ is derived as usual in the computation of successor DBM zones (see, for instance, [41]), and the density-function $f_{Q^b}()$ is obtained by shifting all the components by the value of $\tau^a(t_o)$ and by eliminating $\tau^a(t_o)$ through a projection:

$$\begin{aligned} &f_{Q^b}(x_1, \ldots, x_{I-1}) \\ &\quad = \int_{E_o^+(x_1,\ldots,x_{I-1})}^{L_o^+(x_1,\ldots,x_{I-1})} f_{Q^a}(x_o, x_1 + x_o, \ldots, x_{I-1} + x_o) dx_o, \end{aligned} \tag{27}$$

where $I = |Q^a|$ and where $E_o^+(x_1, \ldots, x_{I-1})$ and $L_o^+(x_1, \ldots, x_{I-1})$ denote the minimum and the maximum values of $x_o$ such that $\langle x_o, x_1 + x_o, \ldots, x_{I-1} + x_o\rangle \in D_{Q^a}$, respectively.

If $U^a \neq \phi$, denote by $t_j$ the deterministic transition that becomes distributed and is added to $Q^b$, and assume that $Q^b$ is encoded in a vector putting in the first $I - 1$ positions the same elements that were in $Q^a$ in the same order, and that $t_j$ is encoded in the last position. If we denote by $\delta()$ the Dirac

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

Delta function, the density-function of distributed transitions can be written as

$$\begin{aligned} & f_{Q^b}(x_1,\ldots,x_{I-1},x_I) \\ & \quad = \int_{E_o^+(x_1,\ldots,x_{I-1},x_I)}^{L_o^+(x_1,\ldots,x_{I-1},x_I)} f_{Q^a}(x_o,x_1+x_o,\ldots,x_{I-1}+x_o) \\ & \qquad\qquad \cdot \delta(x_I+x_o-B_{j*})dx_o \\ & \quad = f_{Q^a}(-x_I+B_{u'_j*},x_1-x_I+B_{u'_j*},\ldots,x_{I-1}-x_I+B_{u'_j*}). \end{aligned} \tag{28}$$

It is worth noting the major difference between the cases of (28) and (27): If there is at least one deterministic timer, then one of the deterministic timers becomes distributed and "takes the memory" of the distribution of $\tau(t_o)$ so that $f_{Q^b}()$ is expressed as a function of $f_{Q^a}()$ independently, whereas, if there is no deterministic timer, then the projection yields an integration that marginalizes $f_{Q^a}()$ with respect to $\tau(t_o)$ (see (27)). In the latter case, the resulting density-function $f_{Q^b}()$ depends on the function that expresses integration extrema $E_o^+()$ and $L_o^+()$ (see (28)), coupling the calculus of density-functions and DBM supports [29].

#### 4.3.3 Disabling

When $t_o$ fires one or more transitions may be disabled. This yields a new vector of times to fire $\tau^c$, derived by repeatedly eliminating the time-to-fire of each disabled transition $t_d$. If $\tau^b(t_d)$ is deterministic or synchronized, it can be removed without changing the set of distributed transitions and thus leaving unchanged their joint density-function so that $f_{Q^c}(x) = f_{Q^b}(x)$. The same happens if $\tau^b(t_d)$ is distributed but it is synchronized with some $\tau^b(t_\delta) \in V^b$, as this case can be reduced to the previous one by swapping $\tau^b(t_d)$ and $\tau^b(t_\delta)$ in the partition between the sets $Q^b$ and $V^b$.

Whereas, if $\tau^b(t_d)$ is distributed and it is not synchronized with any element of $V^b$, then its elimination reduces by one element the set of distributed transitions: $Q^c = Q^b \backslash \{\tau^b(t_d)\}$. The density-function $f_{Q^c}()$ is thus obtained through a projection operation that marginalizes $f_{Q^b}()$ with respect to $\tau^b(t_d)$. If we assume that $Q^b$ is encoded in a vector where $\tau^b(t_d)$ occupies the first position, the derivation can be expressed as

$$\begin{aligned} & f_{Q^c}(x_1,\ldots,x_{I_b-1}) \\ & \quad = \int_{E_d(x_1,\ldots,x_{I_b-1})}^{L_d(x_1,\ldots,x_{I_b-1})} f_{Q^b}(x_d,x_1,\ldots,x_{I_b-1})dx_d, \end{aligned} \tag{29}$$

where $I_b = |Q^b|$ and where $E_d(x_1,\ldots,x_{I_b-1})$ and $L_d(x_1,\ldots,x_{I_b-1})$ denote the minimum and the maximum values of $x_d$ such that $\langle x_d,x_1,\ldots,x_{I_b-1}\rangle \in D_{Q^b}$, respectively.

#### 4.3.4 Newly Enabling

The vector $\tau'$ of times to fire of transitions enabled in the child class $\Sigma' = \langle m', D', f_{Q'}()\rangle$ is finally obtained by extending $\tau^c$ with the times to fire of transitions newly enabled in $\Sigma'$, each distributed according to its own static density-function. Newly enabled transitions with point-like static firing intervals (i.e., $EFT^s = LFT^s$) are added to $U'$, while those with a distributed static firing interval (i.e., $EFT^s < LFT^s$) are added to $Q'$. According to the semantics of sTPNs, their time-to-fire is sampled independently from the times to fire of persistent transitions, and their static density-function is thus added in the form of a product.

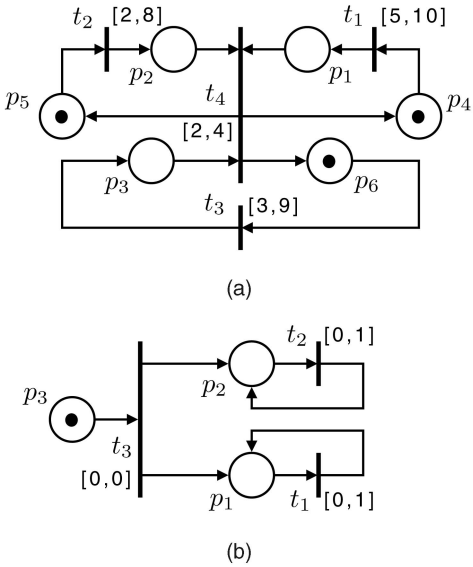


Fig. 3. Two stochastic Time Petri Nets. Static distributions of nondeterministic transitions are assumed uniform.

To express the derivation, we assume that $Q'$ is encoded in a vector where newly enabled transitions take the last positions, with indexes ranging from $I^c$ to $I'-1$. For simplicity, we also denote by $t_i$ the newly enabled transition whose time-to-fire takes position $i$ in the vector encoding of $Q'$. Under these notational assumptions:

$$f_{Q'}(x_o,\ldots,x_{I^c-1},\ldots,x_{I'-1}) = f_{Q^c}(x_o,\ldots,x_{I^c-1}) \cdot \prod_{i=I^c}^{I'-1} f_{t_i}^s(x_i). \tag{30}$$

## 5 Analysis

Enumeration of the stochastic reachability relation $\Sigma^p \overset{t,\mu}{\Longrightarrow} \Sigma^c$ expands the class-graph into a stochastic timed transition system, called here *stochastic-graph*, where vertices are classes labeled with a state-density-function and edges are net transitions labeled with a measure of probability. If the stochastic-graph is finite, its analysis supports derivation of various relevant reward measures.

### 5.1 Boundedness and Regeneration Classes

The stochastic-graph may include multiple stochastic-classes expanding the same class with different state-density-functions. The case is well exemplified with reference to the net of Fig. 3a and its graphs reported in Fig. 4. The class-graph in Fig. 4a includes a "diamond structure" composed of four classes $S^0$, $S^2$, $S^3$, and $S^6$, corresponding to different orderings in the firing of transitions $t_1$ and $t_2$: The order does not influence the set of feasible future behaviors, but it conditions their distribution of probability; according to this, in the stochastic-graph of Fig. 4b, class $S^4$ is expanded in two different stochastic-classes $\Sigma^{4.1}$ and $\Sigma^{4.2}$, which share the same marking and domain but differ in the state-density-function.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

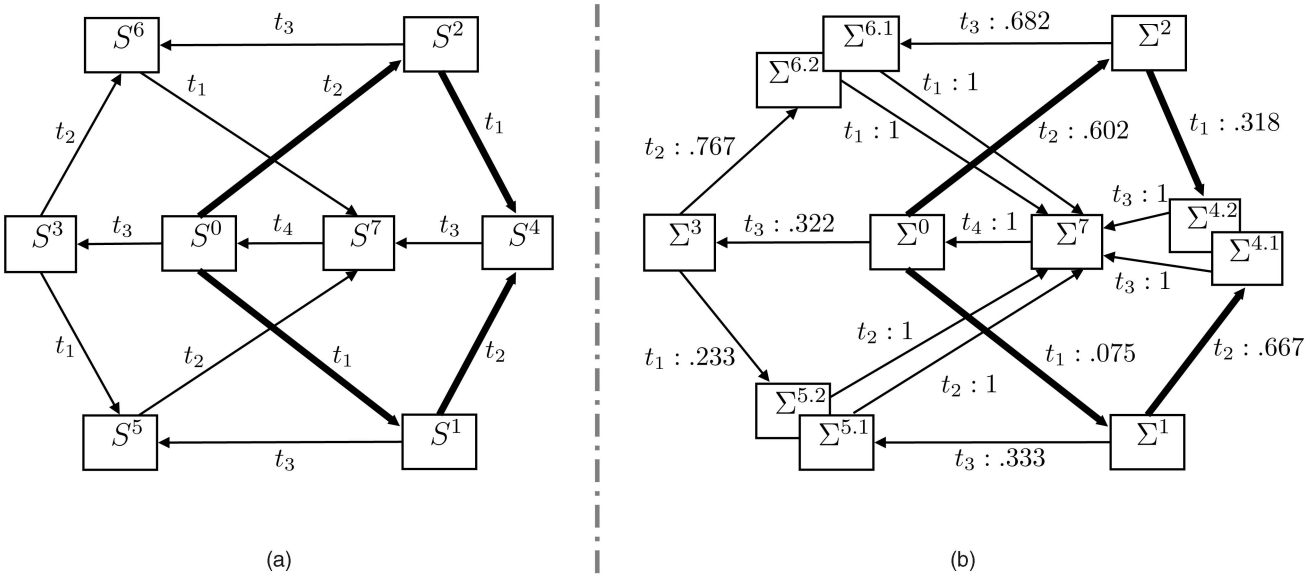


Fig. 4. (a) The class-graph of the net in Fig. 3a. Each vertex is a class $\langle m, D\rangle$. (b) The corresponding stochastic-graph. Each vertex is a stochastic-class $\langle m, D, f_Q()\rangle$ and edges are labeled with a measure of probability. Stochastic classes labeled by x.n share the marking and the time-domain of the class labeled by x, but they differ in the state-density-function $f_Q()$.

Due to the one-to-many relation between classes and stochastic-classes, the stochastic-graph of a model with a finite class-graph can be unbounded. The case is well exemplified by the class-graph of Fig. 5a (derived from the net in Fig. 3b): In the infinite repetition of the firing of $t_1$ from the class $S^2$, transition $t_1$ repeatedly fires and re-enables itself, leaving $t_2$ persistent; the probability that $t_1$ fires before $t_2$ becomes lower at each firing, as $t_1$ always resamples its time-to-fire in the interval $[0, 1]$ with its uniform static density distribution, while $t_2$ accumulates memory of a growing number of events in which it has been overtaken by transition $t_1$; this produces an infinite sequence of different density-functions where $t_2$ becomes more and more concentrated close to 0, as shown in Fig. 5b. Still referring to the class-graph of Fig. 5a, a similar accumulation of memory occurs also along the cycle $S^2 \xrightarrow{t_2} S^1 \xrightarrow{t_1} S^2$ where $t_1$ and $t_2$ fire alternately producing an infinite sequence of density-functions converging to the cosine-shaped form shown in Fig. 5c: In this case, there is not a single transition accumulating memory through persistence, but memory is passed among transition instances with overlapping enabling periods.

The examples of Fig. 5 suggest that the unbounded expansion of a class into an infinite number of stochastic-classes is related to cycles in the class-graph where some memory is passed among persistent transitions across every firing. Developing on this intuition, we introduce a concept of regeneration-class, which breaks memory propagation and guarantees a sufficient condition for the boundedness of the stochastic-graph.

**Definition 5.1.** *A* regeneration-class *is a class where every enabled transition: 1) is newly enabled or 2) has an exponential static density-function with support $[0, \infty)$ or 3) is constrained to assume a deterministic value or a deterministic delay with respect to a transition satisfying any of the two previous conditions.*

A regeneration-class can be associated with a unique state-density-function. In fact, by the semantics of sTPNs, newly enabled transitions are distributed independently according to their static density-function, while values for synchronized or deterministic transitions are determined by the domain of the state-class. Besides, transitions with static density-function exponential over $[0, \infty]$ always maintain their distribution by inductive extension of the following lemma, which basically restates the memoryless property of exponentially distributed times to fire.

**Lemma 5.1.** *Let $\tau(t_e)$ be the time-to-fire of a transition enabled in a class $\Sigma$, and let it be exponentially distributed over $[0, \infty]$ independently from any other transition enabled in $\Sigma$. Let $\Sigma'$ be the successor of $\Sigma$ through the firing of a transition $t_o \neq t_e$, and let $t_e$ be persistent at the firing of $t_o$. In $\Sigma'$, $\tau(t_e)$ is still exponentially distributed and independently from the time-to-fire of every other enabled transition.*

This permits us to demonstrate the following relevant result, which characterizes the conditions for the applicability of the proposed technique in its present formulation.

**Theorem 5.1.** *If $G$ is a finite class-graph where every cyclic path traverses at least one regeneration-class, then the stochastic-graph $\Gamma$ associated with $G$ is also finite.*

Regeneration classes can be identified by inspection of the class-graph, prior to the stochastic expansion. The hypothesis of Theorem 5.1 can thus be easily checked by considering a reduced class-graph $G^-$ derived from $G$ by removing every regeneration-class, and then checking whether $G^-$ includes any cycle. The test runs in linear time with respect to the size of $G$ and does not require the actual construction of $G^-$. Application of the test gives a positive result (i.e., every detected cycle visits at least one regeneration-class) in most examples of TPNs reported in the literature and in particular in those of [7], [43], [28].

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

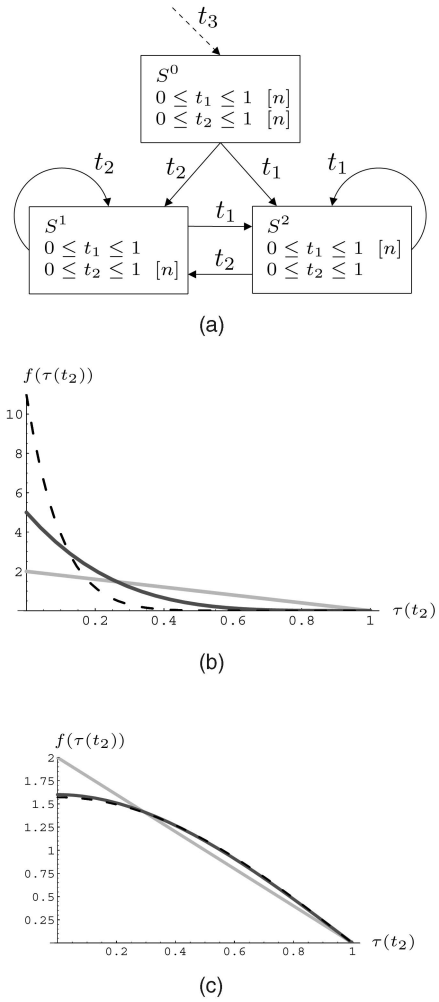


Fig. 5. (a) The class-graph of the net in Fig. 3b. (b) State-density-functions at the first, fourth, and 10th expansion of class $S^2$ in the repeated execution of the self-loop $S^2 \xrightarrow{t_1} S^2$. (c) State-density-functions at the first, third, and 11th expansion of class $S^2$ in the repeated execution of the alternate cycle $S^2 \xrightarrow{t_2} S^1 \xrightarrow{t_1} S^2$.

Conversely, the test fails in the examples reported in [41], [21] (due the overlap between activity cycles of transitions modeling complex arrival processes), in the M/G/1 model of [16] and [24] (due to a possible cycle in the Continuous-Time Markovian Chain (CTMC) subordinated to the activity period of a general transition), in the M/D/2/K model of [31] (due to a possible cycle that alternates the firing of deterministic transitions with an unbounded number of subsequent overlaps).

## 5.2 Probability Distribution of Sequence Timings

The succession relation between stochastic-classes can be inductively extended to establish an isomorphism between the *firing sequences* of an sTPN and the *paths* in its stochastic-graph. In fact, by transitive closure of Definition 3.2: If the sTPN is in a marking $m$ and its timers are a random variable distributed over a domain $D$ with state-density-function $f_Q()$, then a firing sequence $\rho$ can be executed if and only if $\rho$ is a path originating from class $\Sigma^0 = \langle m, D, f_Q() \rangle$ in the stochastic-graph; in this case, the probability that the net executes along $\rho$ is the product of probabilities labeled on

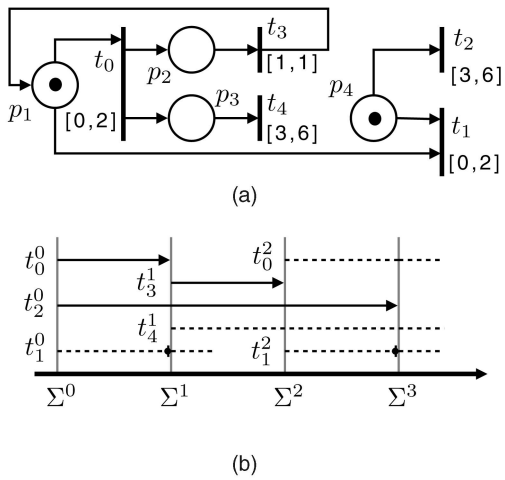


Fig. 6. A simple net and the schema of one of its traces: $\rho = \Sigma^0 \xrightarrow{t_0} \Sigma^1 \xrightarrow{t_3} \Sigma^2 \xrightarrow{t_2} \Sigma^3$. Dotted lines denote times to fire sampled by transitions that do not come to fire; a dot marks the point where the transition is disabled.

the edges traversed by the path itself, the state at the end of the sequence has the marking of the class $\Sigma^N = \langle m^N, D^N, f_Q^N() \rangle$ reached by the path, and its timers are a random variable distributed over $D^N$ with density $f_Q^N()$.

This correspondence enables quantitative evaluation of the relative probability of critical sequences selected in the range of feasible behaviors accepted by a model. This result is now extended with the evaluation of a measure of probability over the space of feasible timings of any single firing sequence.

### 5.2.1 Firing Sequence Notation

We consider a firing sequence $\rho$ that originates from class $\Sigma^0$, visits classes $\Sigma^1$ through $\Sigma^{N-1}$ and terminates in $\Sigma^N$. We assume that $\Sigma^0$ is a regeneration-class. The assumption is in fact not necessary, but it is not restrictive for the theory that will be developed in Section 5.3 and it largely simplifies notational complexity: The assumption guarantees that all distributed times to fire are either associated with newly enabled transitions or they are exponentially distributed over $[0, \infty]$; in both cases, they are all independent and not limited by any cross-constraint.

To give identity to transitions enabled in multiple classes, a transition $t_i$ newly enabled in $\Sigma^n$ will be denoted by $t_i^n$ in $\Sigma^n$ and in all the subsequent classes where it is persistent. In so doing, $t_i^n$ identifies not only a transition in the static structure of the model, but also a *transition activation* and thus a timer in its dynamic execution (e.g., in Fig. 6, $t_2^0$ denotes the activation of transition $t_2$ that is newly enabled in $\Sigma^0$ and persistent until the firing that enters class $\Sigma^3$).

Each transition activation $t_i^n$ is associated with an *absolute virtual firing time* (firing time for short) denoted by $\tau_i^n$ and defined as the sum of the time-to-fire taken by $t_i$ at its newly enabling plus the time elapsed from the start of the sequence $\rho$ to the firing that enters $\Sigma^n$. We say that $\tau_i^n$ is absolute as it is referred to the start-time of the trace, and that it is virtual as $t_i^n$ might not come to fire, either because it is disabled or because the sequence $\rho$ terminates (e.g., in Fig. 6, $t_1^0$ is disabled by $t_0^0$ and $t_0^2$ is persistent when the trace $\rho$ terminates).

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

To identify enabling and disabling relations among transitions, let $\iota(n)$ be index of the transition that enters class $\Sigma^n$ and let $\nu(n)$ be its enabling class, so that $\tau_{\iota(n)}^{\nu(n)}$ is the time at which $\Sigma^n$ is entered (e.g., in Fig. 6, $t_{\iota(3)}^{\nu(3)} = t_2^0$). Finally, for any transition $t_i^n$ disabled along the sequence $\rho$, let $\gamma(i,n)$ be the index of the transition that disables $t_i^n$ and let $\delta(i,n)$ be its enabling class, so that $\tau_{\gamma(i,n)}^{\delta(i,n)}$ denotes the time at which $t_i^n$ is disabled. In a similar manner, if $t_i^n$ is still persistent when $\rho$ terminates, let $t_{\gamma(i,n)}^{\delta(i,n)}$ be the transition that terminates $\rho$ (e.g., in Fig. 6, $t_1^2$ is disabled by $t_2^0$ and thus $t_{\gamma(1,2)}^{\delta(1,2)} = t_2^0$; and, $t_4^1$ is persistent when $t_2^0$ fires and the trace $\rho$ terminates, and thus $t_{\gamma(4,1)}^{\delta(4,1)} = t_2^0$).

### 5.2.2 Timing Constraints

Virtual firing times of transitions enabled along $\rho$ are subject to three different kinds of constraints:

1. For any transition $t_i^n$, the time-to-fire taken at the newly enabling falls in the interval $[EFT^s(t_i), LFT^s(t_i)]$:

$$\begin{aligned} \forall t_i^n \text{ enabled along } \rho : EFT^s(t_i) &\le \tau_i^n - \tau_{\iota(n)}^{\nu(n)} \\ &\le LFT^s(t_i); \end{aligned} \tag{31}$$

2. the virtual firing time of any transition $t_x^n$ that does not come to fire is not lower than the virtual firing time of the transition that disables $t_x^n$ itself or terminates the sequence $\rho$:

$$\forall t_x^n \text{ enabled but not fired along } \rho : \tau_x^n \ge \tau_{\gamma(x,n)}^{\delta(x,n)}; \tag{32}$$

3. class $\Sigma^n$ is entered not later than $\Sigma^{n+1}$:

$$\forall n \in [0, N-1] : \tau_{\iota(n+1)}^{\nu(n+1)} \ge \tau_{\iota(n)}^{\nu(n)}. \tag{33}$$

Equations (31)-(33) define a domain in DBM form, which constrains the range of variation for the vector of virtual firing times of transitions enabled along $\rho$. Quite intuitively, each of these constraints is sound: any feasible timing for $\rho$ satisfies (31) and (32) as a direct consequence of the semantics of sTPNs, and it satisfies (33) as a consequence of the ordering of events along $\rho$. The following lemma proves that the overall set of constraints is also complete.

**Lemma 5.2.** *Let $\bar{\tau}$ be a vector of virtual firing times for $\rho$ that satisfies (31)-(33). $\bar{\tau}$ determines the value sampled by every transition newly enabled along $\rho$ and thus defines a firing sequence $R$. This sequence is feasible for the sTPN underlying the stochastic-graph.*

The DBM defined by (31)-(33) may have a null measure due to the presence of deterministic or deterministically synchronized virtual firing times: Any transition with a non-point-like static firing interval will have a distributed virtual firing time, provided that $\rho$ has a nonnull probability (i.e., that it does not traverse any edge with null probability), whereas the virtual firing time of any transition with a point-like firing interval is bound to have a deterministic delay with respect to the virtual firing time of the transition that enters its enabling class. The derivation of a density-function for the range of feasible timings of $\rho$ thus requires that $D_\rho$ be represented in partitioned form, as discussed in Section 3.3:

$$\begin{aligned} Q &= \left\{\tau_i^n \,.\, EFT^s(t_i) \neq LFT^s(t_i)\right\}, \\ U &= \{\tau_*\} \bigcup \left\{\tau_h^n \,.\, (EFT^s(t_h) = LFT^s(t_h)) \wedge \left(\tau_{\iota(n)}^{\nu(n)} \in U\right)\right\}, \\ V &= \left\{\tau_k^n \,.\, (EFT^s(t_k) = LFT^s(t_k)) \wedge \left(\tau_{\iota(n)}^{\nu(n)} \in V \bigcup Q\right)\right\}, \end{aligned} \tag{34}$$

where $\tau_* = 0$ denotes the start-time of $\rho$, i.e., the entering time $\tau_{\iota(0)}^{\nu(0)}$ of the initial class $\Sigma^0$.

In the following, we denote by $\tau_{\kappa(n)}^{\mu(n)}$ the distributed transition synchronized with the entering time $\tau_{\iota(n)}^{\nu(n)}$ of class $\Sigma^n$ and we denote by $c^n$ the synchronization delay (i.e., $\tau_{\iota(n)}^{\nu(n)} = \tau_{\kappa(n)}^{\mu(n)} + c^n$ with $\tau_{\kappa(n)}^{\mu(n)} \in Q$); in the special case that $\tau_{\iota(n)}^{\nu(n)}$ is distributed, then $\tau_{\kappa(n)}^{\mu(n)} = \tau_{\iota(n)}^{\nu(n)}$ and $c^n = 0$. We then encode the virtual firing times $\tau_i^n$ of distributed transitions enabled along $\rho$ in a vector $\tau$. We denote by $L$ its length and we assume that $l(i,n)$ is the index of the vector element encoding $\tau_i^n$ (i.e., $\tau_{l(i,n)} = \tau_i^n$).

### 5.2.3 Stochastic Characterization

The stochastic characterization of the vector $\tau$ is defined by the probability distributions according to which newly enabled transitions sample their times to fire and by the conditioning induced by sequencing constraints that let the net run along $\rho$. To separate the two aspects, we introduce a set of auxiliary variables, denoted by $\hat{\tau}_i^n$, that neglect the conditioning imposed by the sequencing constraints. We encode the variables $\hat{\tau}_i^n$ in a vector $\hat{\tau}$ with the same indexing function $l(i,n)$ used for $\tau$, and we denote by $f_{\hat{\tau}}()$ its probability density-function. By construction, the vector $\tau$ can be obtained by conditioning $\hat{\tau}$ to range within $D_\tau$ (i.e., $\tau = \hat{\tau} | \hat{\tau} \in D_\tau$) and thus:

$$f_\tau(x) = \frac{f_{\hat{\tau}}(x)}{\int_{D_\tau} f_{\hat{\tau}}(y)dy}. \tag{35}$$

To derive $f_{\hat{\tau}}(x)$, we consider a further set of variables $\theta$:

$$\theta_i^n = \hat{\tau}_i^n - \hat{\tau}_{\kappa(n)}^{\mu(n)} - c^n, \tag{36}$$

which we encode in a vector $\theta$ according to the same index $l(i,n)$. According to (36), the transformation from $\hat{\tau}$ to $\theta$ can be expressed in matrix form as

$$\theta = A\hat{\tau} + C, \tag{37}$$

where $C$ is a vector of $L$ constants and $A$ is an $L \times L$ array with $det(A) = 1$.

By construction, $\theta_i^n$ represents the time-to-fire sampled by $t_i$ at its newly enabling in class $\Sigma^n$, independently from sequencing constraints imposed by $\rho$. According to this, $\theta_i^n$ is distributed with probability density-function $f_{t_i}^s()$ and the joint probability of vector $\theta$ can be expressed as a product:

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

$$f_\theta(x) = \prod_{t_i^n \in dis(\rho)} f_{t_i}^s\big(x_{l(i,n)}\big), \tag{38}$$

where $dis(\rho)$ denotes the set of distributed transitions enabled along $\rho$.

By replacing variables according to (37), we can thus write:

$$\begin{aligned} f_{\hat{\tau}}(x) &= |det(J)| \cdot f_\theta(\theta(x)) \\ &= |det(J)| \cdot \prod_{t_i^n \in dis(\rho)} f_{t_i}^s\big(x_{l(i,n)} - x_{l(\kappa(n),\mu(n))} - c^n\big), \end{aligned} \tag{39}$$

where $det(J)$ is the determinant of the Jacobian matrix of the transformation $\theta = \theta(\hat{\tau})$. According to (37), $J = A$ and thus $det(J) = 1$. By replacing (39) in (35) we thus finally obtain:

$$f_\tau(x) = \frac{\prod_{t_i^n \in dis(\rho)} f_{t_i}^s\big(x_{l(i,n)} - x_{l(\kappa(n),\mu(n))} - c^n\big)}{\int_{D_\tau} \prod_{t_i^n \in dis(\rho)} f_{t_i}^s\big(y_{l(i,n)} - y_{l(\kappa(n),\mu(n))} - c^n\big) dy}. \tag{40}$$

The density-function $f_\tau()$ enables a straightforward reconstruction of the probability distribution for the vector of virtual firing times of all the transitions enabled along the trace. In turn, this probability distribution provides a measure of probability for the dense variety of timings that can occur in the transient behavior along a feasible firing sequence. This has direct application in a number of relevant quantitative problems, such as the evaluation of the probability that a firing sequence is completed within a hard/firm deadline or the evaluation of the average delay with respect to a soft deadline [22]. In particular, it supports the derivation of the average time spent in the execution of a firing sequence. This provides an instrumental step for steady-state analysis based on Markov Renewal Theory developed in the next section.

### 5.3 Characterization and Analysis of the Underlying Stochastic Process

The stochastic-graph provides a discrete partition of the state-space of an sTPN, abstracting its behavior into a discrete-state continuous-time process $\mathbf{X} = \{X(t) = \Sigma,\ t \in \mathbb{R}\}$ (with $X(t) = \Sigma$ meaning that at time $t$ the net is in the stochastic-class $\Sigma$).

The embedded chain of $\mathbf{X}$, denoted by $\mathbf{Y} = \{Y_n = \Sigma,\ n \in \mathbb{N}\}$ is a discrete time Markov chain (DTMC) as a direct consequence of Definition 3.2. Its analysis yields the transient probability $\eta_{\Sigma^i}^{\mathbf{Y}}(n) = Prob\{Y_n = \Sigma^i\}$ that, after $n$ firings, the net is in a stochastic-class $\Sigma^i$ and the steady-state limit $\eta_{\Sigma^i}^{\mathbf{Y}} = \lim_{n\to\infty} Prob\{Y_n = \Sigma^i\}$ (provided that the chain is irreducible and aperiodic). In turn, transient and steady-state probability of any marking $m$ can be derived by marginalizing measures over the set of classes with marking $m$.

Unfortunately, the lack of any restriction on the number of concurrently enabled nonexponential transitions prevents a straightforward approach based on Markov Renewal Theory to the temporal reconstruction of $\mathbf{X}$ from steady-state probabilities of $\mathbf{Y}$, as sojourn times in stochastic-classes may be nonexponential and mutually dependent (both in the support and in the distribution). To circumvent the limitation we consider a restricted process $\bar{\mathbf{X}}$ derived from $\mathbf{X}$ by observing only regeneration-classes: $\bar{\mathbf{X}} = \{\bar{X}(t) = \Sigma, t \in \mathbb{R}\}$ (with $\bar{X}(t) = \Sigma$ meaning that, at time $t$, $\Sigma$ is the last visited regeneration-class). The states of $\bar{\mathbf{X}}$ are all and only the regeneration-classes of $\mathbf{X}$, and any two states $\Sigma^i$ and $\Sigma^j$ are connected through a transition with probability $\mu_{ij}$ iff $\mathbf{X}$ contains a path from $\Sigma^i$ to $\Sigma^j$ that does not visit any regeneration-class and $\mu_{ij}$ is the probability that $\Sigma^j$ is the first regeneration point reached after leaving $\Sigma^i$. Due to the property of regeneration-classes, $\bar{\mathbf{X}}$ is a continuous-time semi-Markov chain, and its embedded chain $\bar{\mathbf{Y}}$ is thus a discrete time Markov chain.

As a fortunate fact, the same condition that guarantees bounded expansion of the stochastic-graph (i.e., the absence of cyclic paths that do not visit any regeneration-class) also guarantees that the number of paths between any two regeneration-classes is finite. This enables the derivation of probability and average time spent along each path between any two regeneration-classes.

According to Section 5.2, the probability of each path is computed as the product of probabilities of the transitions traversed across the path itself. This enables the analysis of the subordinated process between any two regeneration-classes providing transition probabilities for the DTMC $\bar{\mathbf{Y}}$. In particular, steady-state analysis of $\bar{\mathbf{Y}}$ determines limit probabilities $\eta_{\Sigma^i}^{\bar{\mathbf{Y}}} = \lim_{n\to\infty} Prob\{\bar{Y}_n = \Sigma^i\}$, where $\bar{Y}_n = \Sigma^i$ means that, after $n$ firings, $\Sigma^i$ is the last regeneration-class visited by the process $\mathbf{X}$.

Finally, since $\bar{\mathbf{X}}$ is a semi-Markov process, its limit probabilities $\eta_{\Sigma^i}^{\bar{\mathbf{X}}} = \lim_{t\to\infty} Prob\{\bar{X}(t) = \Sigma^i\}$ can be expressed according to Markov Renewal Theory by combining limit probabilities of the embedded chain $\bar{\mathbf{Y}}$ with average sojourn times spent in the states of $\bar{\mathbf{X}}$ itself:

$$\eta_{\Sigma^i}^{\bar{\mathbf{X}}} = \frac{\eta_{\Sigma^i}^{\bar{\mathbf{Y}}} * E\big[SJ_{\Sigma^i}^{\bar{\mathbf{X}}}\big]}{\sum_{\Sigma^h \in \bar{\mathbf{X}}} \eta_{\Sigma^h}^{\bar{\mathbf{Y}}} * E\big[SJ_{\Sigma^h}^{\bar{\mathbf{X}}}\big]}, \tag{41}$$

where $E[SJ_{\Sigma^i}^{\bar{\mathbf{X}}}]$ denotes the expected sojourn time in the state $\Sigma^i$ of the process $\bar{\mathbf{X}}$ (i.e., the average time elapsed from the arrival in the regeneration-class $\Sigma^i$ to the arrival in any subsequent regeneration-class) and can be expressed as

$$E\big[SJ_{\Sigma^i}^{\bar{\mathbf{X}}}\big] = \sum_{\rho \in R_i} \mu_\rho * E[C(\rho)], \tag{42}$$

where $R_i$ is the set of paths reaching any regeneration-class from $\Sigma^i$ and can be determined from inspection of the stochastic-graph; $\mu_\rho$ is the probability that path $\rho$ is followed from class $\Sigma^i$ and can be computed as the product of transition probabilities on traversed edges; $E[C(\rho)]$ is the average time spent in the execution of $\rho$ and be determined according to the treatment of Section 5.2.

## 6 An Example

The theory of sTPN was fully implemented in the Oris tool [41], [14] in two different versions. A first implementation reported in [13], [37], delegates to the Mathematica Kernel [42] the symbolic calculus of integral functions. A much better performing version, reported in [38], implements the entire calculus in a dedicated C++ library, under the assumption that static probability density-functions of nondeterministic transitions are exp-polynomial functions. We illustrate here the kind of results that can be derived

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

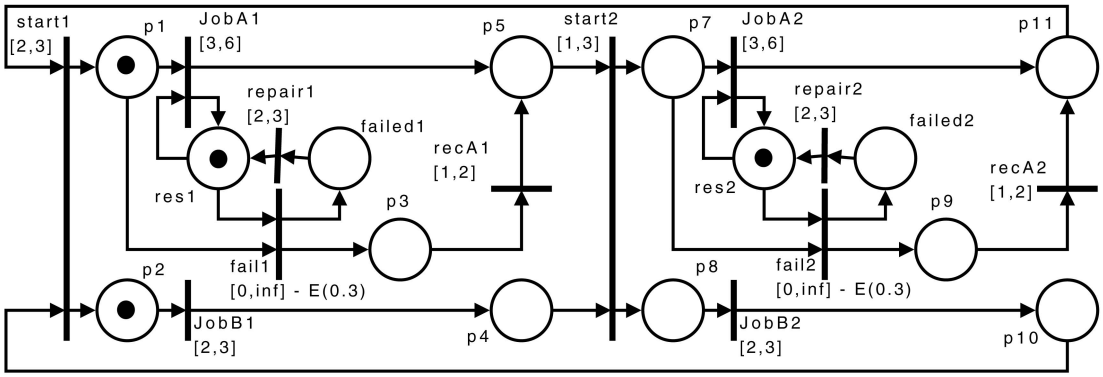


Fig. 7. A stochastic Time Petri Net modeling a system with two production cells. The annotation $\mathrm{E}(\lambda)$ denotes an exponential distribution with rate $\lambda$. For simplicity of presentation, all of the other transitions are supposed to be associated with uniform distributions in their static firing interval. Also, transition weights $\mathcal{C}$ are all made equal to 1 and not reported in the picture.

through the proposed theory and through these implementations, and we point out their present limits of applicability.

### 6.1 The Model

We refer to the sTPN model in Fig. 7. This represents a system composed of two production cells, identified by suffixes 1 and 2, which operate in alternate manner passing control to each other (transitions $start1$ and $start2$). For simplicity of presentation, the two cells have the same structure and timings. As a single intentional exception, transitions $start1$ and $start2$ have different values of the earliest static firing time. Each cell includes two parallel activities named $JobA$ and $JobB$, both associated with uniform duration. In both cells, $JobA$ requires a resource $res$, which may fail during the usage with exponential distribution and rate $\lambda = 0.3$. In case of failure, $JobA$ is not completed, a recovery activity $recA$ with uniform duration is performed, and an action of $repair$ with uniform duration is started.

### 6.2 Nondeterministic Analysis

Nondeterministic analysis produces a graph with 48 state-classes (a fragment is shown in Fig. 8a), which supports verification of timed reachability problems and evaluation of min-max durational properties, through the theory of Time Petri Nets [8], [7], [41] and through the model checker of the Oris tool [37]. For instance, a state-class where both resources are failed is reachable from the initial class through 3 (and only 3) sequences that do not include any firing of $start1$:

$$\begin{aligned}\rho_1 &= S0 : fail1 \rightarrow recA1 \rightarrow JobB1 \rightarrow start2 \rightarrow fail2,\\ \rho_2 &= S0 : fail1 \rightarrow JobB1 \rightarrow recA1 \rightarrow start2 \rightarrow fail2,\\ \rho_3 &= S0 : JobB1 \rightarrow fail1 \rightarrow recA1 \rightarrow start2 \rightarrow fail2.\end{aligned} \quad (43)$$

The duration of three sequences falls in the intervals $[3,5]$, $[3,6]$, and $[4,9]$, respectively, and the maximum feasible

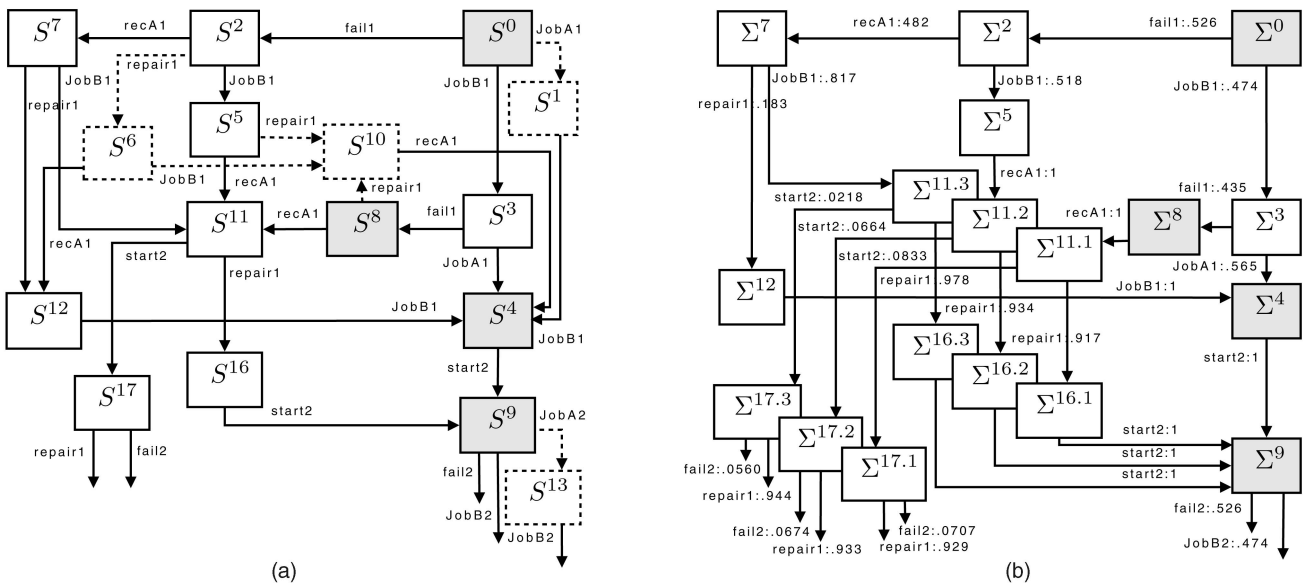


Fig. 8. (a) A fragment of the class-graph, and (b) its stochastic expansion (on the right). $S^0$ is the initial class. Regeneration classes are drawn in gray. Transitions and state-classes that would have null probability are drawn with dotted lines, and they are suppressed in the stochastic expansion. Stochastic classes associated with the same state-class are partially overlapped. Edges between stochastic-classes are annotated with their probability.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

| class | $m$ | $D$ | $f_Q()$ | $\eta^{\bar{\mathbb{Y}}}$ | $E[SJ]$ | $\eta^{\bar{\mathbb{X}}}$ | $E[RT]$ |
|---|---|---|---|---|---|---|---|
| $\Sigma^0$ | p1 res1 p2 res2 | JobA1∈[3,6], fail1∈[0,∞], JobB1∈[2,3] | $0.1e^{-0.3fail1}$ | .333 | 4.45 | .374 | 11.88 |
| $\Sigma^4$ | res1 p4 p5 res2 | start2∈[1,3] | 0.5 | .105 | 2 | .053 | 37.78 |
| $\Sigma^8$ | failed1 p4 res2 p3 | repair1∈[2,3], recA1∈[1,2] | 1 | .069 | 3.88 | .067 | 57.67 |
| $\Sigma^9$ | res1 p7 p8 res2 | JobA2∈[3,6], JobB2∈[2,3], fail2∈[0,∞] | $0.1e^{-0.3fail2}$ | .320 | 4.58 | .370 | 12.38 |
| $\Sigma^{18}$ | res1 p11 p10 res2 | start1∈[2,3] | 1 | .105 | 2.5 | .066 | 37.78 |
| $\Sigma^{19}$ | res1 p10 failed2 p9 | repair2∈[2,3], recA2∈[1,2] | 1 | .069 | 4 | .069 | 57.67 |

Fig. 9. Regeneration classes for the sTPN model of Fig. 7. The last four columns report the results of steady-state analysis of the underlying process: $\eta^{\mathbb{Y}}$ is the limit probability in the embedded DTMC $\bar{\mathbb{Y}}$ of the restricted continuous semi-Markov process $\bar{\mathbb{X}}$ reported in Fig. 13; $E[SJ]$ is the mean sojourn time in $\bar{\mathbb{X}}$; $E[RT]$ is the mean recurrence time derived through Markov Renewal Theory.

sojourn time in the condition of double failure is equal to one time unit for all of them.

The condition of double failure is also reachable through eight sequences that start from the marking $p7\ p8\ res1\ res2$ and do not visit any intermediate firing of $start2$; in this case, the time to reach the condition of double failure falls within the interval [4, 9], and the sojourn time in the state of double failure is bound to be equal to 0.

### 6.3 Stochastic Expansion

Inspection of the class-graph identifies six regeneration-classes, listed in Fig. 9, and shows that every cycle visits at least one regeneration-class. According to Theorem 5.1, the class-graph can thus be expanded into a finite stochastic-graph, which turns out to include 86 stochastic-classes (a fragment is shown in Fig. 8b). As shown in Fig. 8, some state-classes are expanded into multiple stochastic-classes, with the same marking and time-domain but with different state-density-functions (e.g., stochastic-classes $\Sigma^{11.1}$, $\Sigma^{11.2}$, and $\Sigma^{11.3}$, and their respective successors). Conversely, 41 edges among state-classes have null probability and are suppressed in the stochastic expansion (Fig. 8). This also suppresses 23 state-classes (and seven markings) that are reachable only through null probability edges, and a number of critical traces identified in nondeterministic analysis (e.g., all the traces that reach a state with both resources in failed mode starting from the marking $p7\ p8\ res1\ res2$ without any intermediate firing of $start2$).

### 6.4 Stochastic Analysis

The stochastic-graph supports evaluation of various quantitative measures of probability for feasible behaviors identified in nondeterministic analysis. At the first level, probabilities of edges between stochastic-classes provide a measure of the probability for any firing sequence selected in nondeterministic analysis. For instance, for the sequences listed in (43), according to probabilities in Fig. 8b, we have

$$\begin{aligned} P_{\rho_1} &= .526 * .482 * .817 * .0218 * .0560 = 0.000253, \\ P_{\rho_2} &= .526 * .518 * 1 * .0664 * .0674 = 0.00122, \\ P_{\rho_3} &= .474 * .435 * 1 * .0833 * .0707 = 0.00121. \end{aligned} \tag{44}$$

Since different sequences are alternative, the total probability to reach a double failure from the initial class before the next firing of $start1$ is $P_{double-failure} = 0.000253 + 0.00122 + 0.00121$.

At a second level of analysis, the stochastic-graph can be analyzed as a Discrete Time Markov Chain $\mathbb{Y}$ to derive transient and steady-state probabilities of each class and each marking. For instance, Fig. 10 reports steady-state probabilities for classes appearing in the fragment of the stochastic-graph in Fig. 8b.

At a third level of analysis, the technique for trace analysis described in Section 5.2 supports evaluation of the distribution of the time spent in the execution of any feasible firing sequence. This can be combined with the probabilities of individual firing sequences to evaluate the distribution of time spent in the transition between any two stochastic-classes connected through a finite number of paths. Since every cycle in the stochastic-graph visits at least one regeneration-class, the number of paths between any two stochastic-classes is finite, and it is thus possible to evaluate the mean value and the probability density-function for the time spent between any two regeneration-classes. Figs. 11 and 12 illustrate the concept with reference to the case of the two (and only two) firing sequences in the stochastic-graph of Fig. 8 that reach $\Sigma^9$ from $\Sigma^0$ without visiting any intermediate regeneration-class.

The evaluation of the average time spent along each sequence between any two regeneration-classes defines the

| $\Sigma^0$ | $\Sigma^2$ | $\Sigma^3$ | $\Sigma^4$ | $\Sigma^5$ | $\Sigma^7$ | $\Sigma^8$ | $\Sigma^9$ | $\Sigma^{11.1}$ |
|---|---|---|---|---|---|---|---|---|
| .112011 | .0589022 | .0531089 | .0352323 | .0304952 | .0284070 | .0230812 | .107557 | .0232025 |
| $\Sigma^{11.2}$ | $\Sigma^{11.3}$ | $\Sigma^{12}$ | $\Sigma^{16.1}$ | $\Sigma^{16.2}$ | $\Sigma^{16.3}$ | $\Sigma^{17.1}$ | $\Sigma^{17.2}$ | $\Sigma^{17.3}$ |
| .0304952 | .0230812 | .00520451 | .0226964 | .0284704 | .0211577 | .000506101 | .00202473 | .00192343 |

Fig. 10. Steady-state probabilities for the discrete time Markov chain embedded in the stochastic-graph of Fig. 8b.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

| Id | $\rho$ | min-max | $\mu_\rho$ | E[C($\rho$)] |
|---|---|---|---|---|
| $Trace_3$ | $\Sigma^0 : fail1 \rightarrow JobB1 \rightarrow recA1 \rightarrow repair1 \rightarrow start2 \rightarrow \Sigma^9$ | [3,8] | 0.254175 | 5.25878 |
| $Trace_4$ | $\Sigma^0 : fail1 \rightarrow recA1 \rightarrow JobB1 \rightarrow repair1 \rightarrow start2 \rightarrow \Sigma^9$ | [3,6] | 0.202626 | 4.55395 |

Fig. 11. Two firing sequences between stochastic-classes $\Sigma^0$ and $\Sigma^9$ in the stochastic-graph of Fig. 8b. min-max is the interval of duration of the sequence, derived through nondeterministic analysis. $\mu_\rho$ is the probability derived from the probabilities of traversed edges. $\mathrm{E}[\mathrm{C}\rho)]$ is the average time spent in the execution of the sequence, derived from the density-functions of Fig. 12.

continuous-time semi-Markov process $\bar{\mathbf{X}}$ which observes the visits in regeneration-classes. For our example, this process is reported in Fig. 13, and the results of its steady-state analysis are reported in the last four columns of Fig. 9.

### 6.5 Model Variations and Limits of Applicability

The analysis of a few variations of the discussed example supports an intuitive understanding of the current practical and theoretical limits of applicability of the proposed theory.

Changes in the topological structure of the model do not have a specific effect on stochastic analysis. If each cell is expanded to include a third parallel activity (e.g., a third JobC) or to add more sequential activities (e.g., JobA is decomposed in two sequential chunks), complexity grows with the same pattern as that of nondeterministic analysis of Time Petri Nets: The number of interleavings augments and so does the number of state-classes. In the specific aspect of stochastic analysis, traces between regeneration-classes become longer, determining a substantial increase of complexity in the derivation of the density-function of (40).

For a given topological structure, changes in temporal parameters have a different impact whether they affect density distributions or their supports. If static density-functions of transitions are changed in the shape but not in the support, the structures of the state-class-graph and of the stochastic-graph do not change and the complexity of derivation remains basically unchanged.

Whereas, as in nondeterministic analysis of Time Petri Nets, a variation in the static firing intervals (i.e., a variation in the support of static density-functions) may add/remove a implicit sequencing constraints among concurrently enabled transitions. This may reduce/augment the number of state-classes and thus the set of feasible firing sequences, with a substantial impact on the complexity and even on the termination of stochastic expansion.

We illustrate the concept with reference to the static earliest firing time of transitions $start1$ e $start2$ in Fig. 7, which were intentionally set equal to 2 and 1, respectively. Nondeterministic analysis of the model shows that under this setting it is not possible that $repair2$ is persistent at the firing of $start1$. This guarantees that the firing of $start1$ always leads to a regeneration-class. However, if the earliest firing time of $start1$ is reduced to 1, nondeterministic analysis detects a firing sequence that encounters a failure in both cells and never visits any regeneration-class. This results in an infinite stochastic-graph and prevents the proposed analysis approach. Actually, the limitation is in some sense intrinsic in the model more than in the analysis technique. In fact, if the class-graph includes a cycle that never visits a regeneration-class, then the model accepts an infinite firing sequence which never reaches a state satisfying the Markov property. Approximation in the representation of state-density-functions seems to be a promising approach to attack the problem, providing a new conceptual framework with respect to existing techniques based on the approximation of static density-functions of transitions.

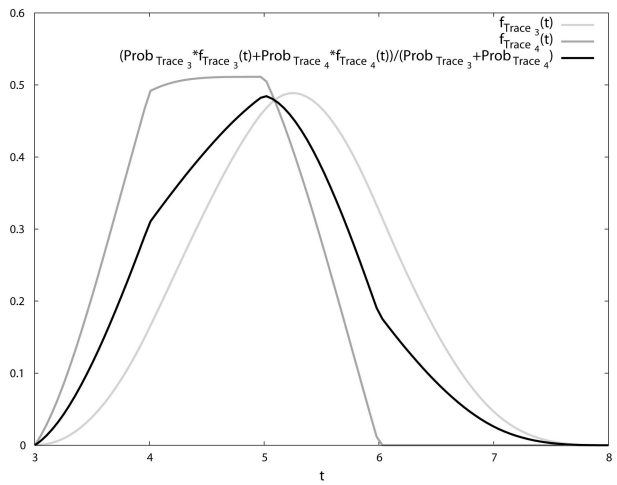


Fig. 12. Probability density-functions for the completion time along the two firing sequences of Fig. 11. The darkest curve is the sum of the two functions weighted according to the probabilities of the two traces.

## 7 Discussion

This paper proposes a new technique for state-space analysis of densely timed models that couples enumeration of DBM state-classes with symbolic derivation of a probability distribution over their support. In so doing, the graph of nondeterministic state-classes [8], [7], [41] is expanded into a stochastic-class-graph that makes explicit transition probabilities among classes and thus enables direct derivation of the relative probability of any symbolic

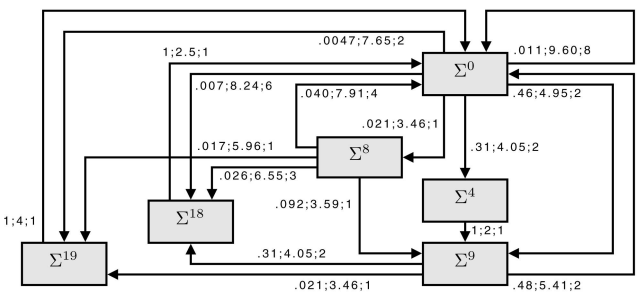


Fig. 13. The continuous-time semi-Markov process $\bar{\mathbf{X}}$ derived from the stochastic-graph of Fig. 8b. Edges are labeled with the probability of the transition, the mean sojourn time in the class in the case of the transition, the number of traces in the stochastic-graph that realize the transition.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

run. The stochastic-class-graph also enables symbolic derivation of the multivariate distribution of sojourn times in the classes visited along any symbolic run between any two regeneration-classes, whose steady-state probabilities can thus be derived through Markov Renewal Theory. In a conceptual perspective, this complements the qualitative identification of feasible behaviors with quantitative probabilities, providing a means to integrate dependability and performance evaluation in correctness verification techniques [41], [7], [30] and tools [6], [14], [9], [20].

In a different perspective, the proposed technique also comprises a new approach for quantitative evaluation of models with multiple concurrent generally distributed timers, and thus compares with the literature of non-Markovian Stochastic Petri Nets [16], [15], [24], [25], [31], [32]. In [24], the logical state is *supplemented* by the vector of ages of enabled general transitions, and its evolution is described through a set of integro-differential state equations. Probability distributions of the ages accumulated in reachable tangible markings are derived through numeric integration, and marking probabilities are derived through marginalization over the age. Though formulated for any number of concurrent generally distributed transitions, practical tractability requires an *enabling restriction* that assumes that at most one generally distributed transition is enabled in any reachable tangible marking [25]. In [16], [15], the same enabling restriction is exploited to guarantee that the underlying stochastic process behaves either as a CTMC or as a CTMC *subordinated* to the activity period of a generally distributed transition. An embedded DTMC is obtained by sampling the process after any firing with no persistent general transition, i.e., at *regeneration points* where the age of general transitions is null. *Transition probabilities* of the DTMC and *conversion factors* needed to reconstruct measures on the continuous-time process can be derived efficiently for the class of models where generally distributed transitions are expolynomial [16]. The limit of the enabling restriction is overcome in [31], [32] under the assumption that generally distributed transitions are deterministic. The process is sampled at equidistant time points yielding a General-State-Space Markov Chain (GSSMC), whose kernel of transition probabilities at subsequent time points is derived through numerical integration. Limit and transient probability distributions for tangible markings supplemented with times to fire of enabled deterministic transitions are thus derived through the formulation of steady-state and transient equations which are solved through numeric integration of a set of Volterra differential equations [31] and of a set of Fredholm integral equations [32], respectively.

Our technique faces the same problem of characterizing the evolution of a stochastic process where memory is carried not only by the discrete-state location but also by a vector of densely valued timers. However, the way how the problem is attacked is basically different in various technical aspects.

A major conceptual difference regards the discrete component of the state to which we associate a continuous probability. While analytical methods associate probability distributions with tangible markings, we distinguish state-classes that have the same marking but a different support for enabled transitions or a different distribution of their joint probability. This refinement permits to associate each stochastic state-class with a multivariate state-density-function supported over a DBM domain, which characterizes the transient behavior until the next firing.

This also permits the derivation of multivariate state-density-functions without numeric integration of differential equations, through a symbolic calculus that accepts a closed-form for the class of models with immediate, deterministic, or expolynomial transitions [38], [29]. In this calculus, the evolution of analytic forms of state-density-functions appears to be tightly intertwined with the algebra of DBM domains, which becomes essential to manage multiple concurrent nondeterministic transitions with possibly bounded supports.

State-density-functions directly determine the matrix of transition probabilities among state-classes. Sojourn times in subsequent classes turn out to be mutually dependent, but this dependency is broken at regeneration-classes, which basically correspond to regeneration points of Markov Renewal theory. Following the approach of [16], [15], we can thus partition the activity of the process in phases that start at the entrance in each regeneration-class, and derive steady-state probabilities of each phase through the evaluation of the relative probability and the average duration of each symbolic run between any two regeneration-classes.

The different principles of operation result in different limits of applicability. Our technique is able to manage models with multiple concurrent nondeterministic generally distributed timers with possibly bounded support. This includes a large class of models (e.g., the net in Fig. 7) that were not tractable with any previous technique, either for structural restrictions [16], [15], [31], [32], or for practical reasons [24]. On the other hand, due to the need to enumerate stochastic states observed after the firing of any transition, our approach is limited to the class of models for which any cycle in the state-space visits at least one regeneration-class (Theorem 5.1). This rules out cases that can be managed in the theory of [31], [32] (e.g., any case including a cycle that alternates the firing of deterministic transitions with an unbounded number of subsequent overlaps), and also cases where the activity cycles of generally distributed transitions do not overlap as required for the application of [16], [15], or [24] (e.g., any case where the Markov chain subordinated to a generally distributed transition includes a cycle).

The comparison of the limitations encountered by the different approaches, and more generally the different perspectives that they provide on the problem of models with (concurrent) generally distributed transitions, suggests directions for further research aimed at combining the analytical techniques of the literature of non-Markovian Petri Nets with the graph-theoretical approach proposed in this paper.

Following a further direction, in [29], we specialize and develop the theory reported in this paper by introducing a mechanism of approximation in the symbolic representation and comparison of multivariate distributions over DBM domains, which smoothes both the computational complexity and the theoretical limitation of unbounded stochastic expansion in the symbolic calculus of stochastic state-classes. In particular, in the theoretical perspective, the

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

approach appears a good basis to attack models with cycles that never visits a regeneration-class.

## APPENDIX: PROOFS

**Lemma 4.1.** *If $\mu_{t_o} > 0$, then the partition $\langle Q, U, V\rangle$ is not changed by the conditioning induced by the assumption that the time-to-fire of $t_o$ is minimum.*

**Proof.** If $Q = \phi$, then $D$ involves only deterministic times to fire and its space of solutions is made of a single point. Since the solutions of $D_{t_o}$ are a subset of those of $D$, the space of solutions of $D_{t_o}$ is also a point, and allof the times to fire in $D_{t_o}$ are thus deterministic.

If $Q \neq \phi$, according to (22), the projection of $D_{t_o}$ over $Q$ has a nonnull measure and thus all the elements of $Q$ are still distributed under the conditioning $\tau(t_o) \leq \tau(t_n)$.

Besides, if $t_k \in V$, then there exists $\tau(t_{\iota_k}) \in Q$ such that the difference $\tau(t_{\iota_k}) - \tau(t_k)$ accepts a single value in the space of solutions of $D$. Since the space of solutions of $D_{t_0}$ is a subset of those of $D$, the difference $\tau(t_{\iota_k}) - \tau(t_k)$ accepts a single solution in $D_{t_o}$. Since $t_{\iota_k}$ remains distributed under the conditioning, $t_k$ remains synchronized.

Finally, if $\tau(t_h) \in U$, then it accepts a single value in the space of solutions of $D$ and also in the space of solutions of $D_{t_o}$ which are a subset of those of $D$. □

**Lemma 5.1.** *Let $\tau(t_e)$ be the time-to-fire of a transition enabled in a class $\Sigma$, and let it be exponentially distributed over $[0, \infty]$ independently from any other transition enabled in $\Sigma$. Let $\Sigma'$ be the successor of $\Sigma$ through the firing of a transition $t_o \neq t_e$, and let $t_e$ be persistent at the firing of $t_o$.*

*In $\Sigma', \tau(t_e)$ is still exponentially distributed and independently from the time-to-fire of every other enabled transition.*

**Proof.** Since $\tau(t_e)$ is independent, the state-density-function of $\Sigma$ can be expressed as a product:

$$f_{\tau(t_e),\tau}(x_e, x) = \lambda e^{-\lambda x_e} f_\tau(x). \tag{45}$$

Following the steps of Section 4.3, it can be easily verified that this form is maintained through the four steps of the successor derivation, provided that integration extrema $E_0^+(), L_0^+(), E_d()$, and $L_d()$ appearing in (27)-(29) do not depend on $x_e$.

In turn, this is guaranteed by the fact that $\tau(t_e)$ is supported over $[0, \infty]$, in the normal form of the time-domain $-b_{*e} = 0$ and $b_{e*} = \infty$. This guarantees that the range of any time-to-fire $\tau(t_i)$ is not upper or lower bounded by any constraint on the difference $\tau(t_i) - \tau(t_e)$.□

**Theorem 5.1.** *If $G$ is a finite class-graph where every cyclic path traverses at least one regeneration-class, then the stochastic-graph $\Gamma$ associated with $G$ is also finite.*

**Proof.** Ab absurdo, let $\Gamma$ be unbounded. Since each stochastic-class $\Sigma \in \Gamma$ is associated with a class $S \in G$, there exists a class $S^o \in G$ that is associated with an unbounded number of stochastic-classes.

This implies that the class-graph $G$ includes a cyclic path $r$ originating in $S^o$, and that the stochastic state graph includes a stochastic-class $\Sigma_k^0$ associated with $S^o$, such that, if $\rho_k$ is the path in the stochastic-graph corresponding to $r$ and originating from $\Sigma_k^0$, then the stochastic-class $\Sigma_{k+1}^0$ reached from $\Sigma_k^0$ through the path $\rho_k$ is other than $\Sigma_k^0$:

$$\begin{aligned} r &= S^o \xrightarrow{t_o} S^1 \xrightarrow{t_1} \cdots \xrightarrow{t_{N-1}} S^o, \\ \rho_k &= \Sigma_k^0 \overset{t_o,\mu_k^o}{\Longrightarrow} \Sigma_k^1 \overset{t_1,\mu_k^1}{\Longrightarrow} \cdots \overset{t_{N-1},\mu_k^{N-1}}{\Longrightarrow} \Sigma_{k+1}^0, \\ &\text{with } \Sigma_k^0 \neq \Sigma_{k+1}^N. \end{aligned} \tag{46}$$

Equation (46) can be easily extended to show that the path $\rho_{k+1}$ that follows the transitions of $r$ starting from the stochastic-class $\Sigma_{k+1}^0$ visits a sequence of stochastic-classes that are all other than the corresponding classes visited along the path $\rho_k$:

$$\begin{aligned} \rho_{k+1} = \Sigma_{k+1}^0 &\overset{t_o,\mu_{k+1}^o}{\Longrightarrow} \Sigma_{k+1}^1 \overset{t_1,\mu_{k+1}^1}{\Longrightarrow} \cdots \overset{t_{N-1},\mu_{k+1}^{N-1}}{\Longrightarrow} \Sigma_{k+2}^0, \\ &\text{with } \Sigma_k^n \neq \Sigma_{k+1}^n \ \forall n = 0, N-1. \end{aligned} \tag{47}$$

Since $r$ is a cyclic path, it visits a regeneration-class, denoted by $S^{n_l}$.

Since $S^{n_l}$ is visited along $r$, it is also associated with two stochastic-classes $\Sigma_k^{n_l}$ and $\Sigma_{k+1}^{n_l}$ visited along $\rho_k$ and $\rho_{k+1}$, respectively. According to (47), $\Sigma_k^{n_l}$ must be other than $\Sigma_{k+1}^{n_l}$, which is not possible as $\Sigma_k^{n_l}$ and $\Sigma_{k+1}^{n_l}$ are stochastic-classes associated with the same regeneration-class. □

**Lemma 5.2.** *Let $\bar{\tau}$ be a vector of virtual firing times for $\rho$ that satisfies (31)-(33). $\bar{\tau}$ determines the value sampled by every transition newly enabled along $\rho$ and thus defines a firing sequence $R$. This sequence is feasible for the sTPN underlying the stochastic-graph.*

**Proof.** By induction, we assume that the sTPN has executed a prefix of $k$ transitions along $R$ following the sequencing of $\rho$ and sampling times to fire of newly enabled transitions according to $\bar{\tau}$. We call $s^k$ the reached state, and we show that transitions newly enabled in $s^k$ can sample the time-to-fire prescribed by $\bar{\tau}$; the $(k+1)$th transition of $\rho$ is firable in $s^k$.

Since $\rho$ is a path in the stochastic-graph, it can be easily shown that transitions newly enabled in $s^k$ are all and only those for which $\bar{\tau}$ prescribes a value of the time-to-fire. According to (31), for each newly enabled transition $t_i^k$, the time-to-fire prescribed by $\bar{\tau}$ falls in the static firing interval of $t_i$ and it is thus feasible.

Besides, the assumption that $\rho$ is a path in the stochastic-graph also implies that the $(k+1)$th transition of $\rho$, denoted by $t_k^m$, is enabled in $s^k$. To conclude the inductive step, it is thus sufficient to prove that $t_k^m$ is also firable, i.e., that its time-to-fire is minimum among all transitions enabled in $s^k$. Ab absurdo, suppose that $t_j^p$ is enabled in $s^k$ and $\tau_k^m > \tau_j^p$. If $t_j^p$ is one of the transitions that are supposed to fire along $\rho$, then the transitive closure of the constraints of (33) implies that $\tau_k^m \leq \tau_j^p$. Whereas, if $t_j^p$ is not supposed to fire along $\rho$, then, according to (32), $\tau_j^p \geq \tau_{\gamma(j,p)}^{\delta(j,p)}$. In turn, the transitive closure of (33) implies that $\tau_k^m \leq \tau_{\gamma(j,p)}^{\delta(j,p)}$. □

## ACKNOWLEDGMENTS

This work was partially supported by the project WiSeDeMon funded by the Italian Ministry of University and Research as a part of the PRIN 2007 Program.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

## 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. Computer Systems,* vol. 2, pp. 93-122, 1984.

[2] M.M. Ajmone and G. Chiola, "On Petri Nets with Deterministic and Exponentially Distributed Firing Times," *Lecture Notes in Computer Science,* vol. 266, pp. 132-145, 1987.

[3] R. Alur, C. Courcoubetis, and D. Dill, "Model-Checking for Real-Time Systems," *Proc. Fifth Symp. Logic in Computer Science,* June 1990.

[4] R. Alur and D.L. Dill, "Automata for Modeling Real-Time Systems," *Proc. 17th Int'l Colloquium on Automata, Languages, and Programming,* 1990.

[5] R. Alur, T.A. Henzinger, and P.H. Ho, "Automatic Symbolic Verification of Embedded Systems," *IEEE Trans. Software Eng.,* vol. 22, no. 3, pp. 181-201, Mar. 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,* Springer, 1996.

[7] B. Berthomieu and M. Diaz, "Modeling and Verification of Time Dependent Systems Using Time Petri Nets," *IEEE Trans. Software Eng.,* vol. 17, no. 3, pp. 259-273, Mar. 1991.

[8] B. Berthomieu and M. Menasche, "An Enumerative Approach for Analyzing Time Petri Nets," *Proc. Int'l Federation for Information Processing Congress 1983,* R.E.A. Mason, ed., vol. 9, pp. 41-46, 1983.

[9] B. Berthomieu, P.-O. Ribet, and F. Vernadat, "The Tool TINA—Construction of Abstract State Spaces for Petri Nets and Time Petri Nets," *Int'l J. Production Research,* vol. 42, no. 14, pp. 2741-2756, 2004.

[10] A. Bobbio, A. Puliafito, and M. Telek, "A Modelling Framework to Implement Preemption Policies in Non-Markovian SPNs," *IEEE Trans. Software Eng.,* vol. 26, no. 1, pp. 36-54, Jan. 2000.

[11] A. Bobbio and M. Telek, "A Benchmark for PH Estimation Algorithms: Result for Acyclic-PH," *Stochastic Models,* vol. 10, pp. 661-677, 1994.

[12] A. Bobbio and M. Telek, "Markov Regenerative SPN with Non-Overlapping Activity Cycles," *Proc. Int'l Computer Performance and Dependability Symp.,* pp. 124-133, 1995.

[13] G. Bucci, R. Piovosi, L. Sassoli, and E. Vicario, "Introducing Probability within State Class Analysis of Dense Time Dependent Systems," *Proc. Second Int'l Conf. Quantitative Evaluation of Systems,* Sept. 2005.

[14] G. Bucci, L. Sassoli, and E. Vicario, "Oris: A Tool for State Space Analysis of Real-Time Preemptive Systems," *Proc. First Int'l Conf. Quantitative Evaluation of Systems,* 2004.

[15] H. Choi, V.G. Kulkarni, and K. Trivedi, "Markov Regenerative Stochastic Petri Nets," *Performance Evaluation,* vol. 20, pp. 337-357, 1994.

[16] G. Ciardo, R. German, and C. Lindemann, "A Characterization of the Stochastic Process Underlying a Stochastic Petri Net," *IEEE Trans. Software Eng.,* vol. 20, no. 7, pp. 506-515, July 1994.

[17] G. Ciardo, J. Muppala, and K.S. Trivedi, "SPNP: Stochastic Petri Net Package," *Proc. IEEE Int'l Workshop Petri Nets and Performance Models,* pp. 142-151, 1989.

[18] C. Daws, A. Olivero, S. Tripakis, and S. Yovine, "The Tool KRONOS," *Hybrid Systems III,* Springer, 1995.

[19] D. Dill, "Timing Assumptions and Verification of Finite-State Concurrent Systems," *Proc. Workshop Computer Aided Verification Methods for Finite State Systems,* 1989.

[20] G. Gardey, D. Lime, M. Magnin, and O. Roux, "Roméo: A Tool for Analyzing Time Petri Nets," *Proc. 17th Int'l Conf. Computer Aided Verification,* 2005.

[21] G. Bucci, A. Fedeli, L. Sassoli, and E. Vicario, "Timed State Space Analysis of Real Time Preemptive Systems," *IEEE Trans. Software Eng.,* vol. 30, no. 2 pp. 97-111, Feb. 2004.

[22] G. Buttazzo, *Hard Real-Time Computing Systems.* Springer, 2005.

[23] R. German, *Performance Analysis of Communication Systems with Non-Markovian Stochastic Petri Nets.* John Wiley & Sons, 2000.

[24] R. German and C. Lindemann, "Analysis of Stochastic Petri Nets by the Method of Supplementary Variables," *Performance Evaluation,* vol. 20, pp. 317-335, 1994.

[25] R. German and M. Telek, "Formal Relation of Markov Renewal Theory and Supplementary Variables in the Analysis of Stochastic Petri Nets," *Proc. Eighth Int'l Workshop Petri Nets and Performance Models,* pp. 64-73, 1999.

[26] D. Harel and A. Pnueli, "On the Development of Reactive Systems," *Logics and Models of Concurrent Systems,* Springer, 1985.

[27] ISO/IEC, International Standard 9126:2001, Information Technology—SW Product Quality—Part 1: Quality Model, 2001.

[28] 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. Systems, Man, and Cybernetics,* vol. 29, no. 4, pp. 541-545, Aug. 1999.

[29] L. Carnevali, L. Grassi, and E. Vicario, "State-Density Functions over DBM Domains in the Analysis of Non-Markovian Models," *IEEE Trans. Software Eng.,* vol. 35, no. 2, pp. 178-194, Mar./Apr. 2009.

[30] D. Lime and O.H. Roux, "Expressiveness and Analysis of Scheduling Extended Time Petri Nets," *Proc. Fifth Int'l Federation of Automatic Control Conf. Fieldbus and Their Applications,* 2003.

[31] C. Lindemann and G.S. Schedler, "Numerical Analysis of Deterministic and Stochastic Petri Nets with Concurrent Deterministic Transitions," *Performance Evaluation,* vols. 27/28, pp. 565-582, 1996.

[32] C. Lindemann and A. Thuemmler, "Transient Analysis of Deterministic and Stochastic Petri Nets with Concurrent Deterministic Transitions," *Performance Evaluation,* vols. 36/37, pp. 35-54, 1999.

[33] M.K. Molloy, "Discrete Time Stochastic Petri Nets," *IEEE Trans. Software Eng.,* vol. 11, no. 4, pp. 417-423, Apr. 1985.

[34] M.F. Neuts, *Matrix Geometric Solutions in Stochastic Models.* Johns Hopkins Univ. Press, 1981.

[35] W. Penczek and A. Polrola, "Specification and Model Checking of Temporal Properties in Time Petri Nets and Timed Automata," *Proc. 25th Int'l Conf. Application and Theory of Petri Nets,* June 2004.

[36] K.G. Popstojanova and K.S. Trivedi, "Stochastic Modeling Formalisms for Dependability, Performance and Performability," *Performance Evaluation—Origins and Directions,* pp. 385-404, Springer, 2000.

[37] L. Sassoli and E. Vicario, "Analysis of Real Time Systems through the Oris tool," *Proc. Fourth Int'l Conf. Quantitative Evaluation of Systems,* Sept. 2006.

[38] L. Sassoli and E. Vicario, "Close Form Derivation of State-Density Functions over DBM Domains in the Analysis of Non-Markovian Models," *Proc. Fifth Int'l Conf. Quantitative Evaluation of Systems,* Sept. 2007.

[39] J.A. Stankovic and K. Ramamritham, "What Is Predictability for Real Time Systems," *J. Real Time Systems,* vol. 2, pp. 247-254, 1990.

[40] M. Telek and A. Horvarth, "Supplementary Variable Approach Applied to the Transient Analysis of Age-MRSPNs," *Proc. Int'l Performance and Dependability Symp.,* pp. 44-51, 1998.

[41] E. Vicario, "Static Analysis and Dynamic Steering of Time Dependent Systems Using Time Petri Nets," *IEEE Trans. Software Eng.,* vol. 27, no. 8, pp. 728-748, Aug. 2001.

[42] WolframResearch "Mathemathica 5.2," www.wolfram.com, 2009.

[43] D. Xu, X. He, and Y. Deng, "Compositional Schedulability Analysis of Real-Time Systems Using Time Petri Nets," *IEEE Trans. Software Eng.,* vol. 28, no. 10, pp. 984-996, Oct. 2002.

**Enrico Vicario** received the doctoral degree in electronics engineering and the PhD degree in information and telecommunications engineering from the University of Florence in 1990 and 1994, respectively. He is a full professor of computer science at the School of Engineering of the University of Florence, where he teaches computer programming, software engineering, and software verification and testing. His present research activity is mainly focused on state-space analysis of concurrent timed systems for correctness verification, quantitative evaluation, and real-time testing. He also carries out technological experimentation and transfer in the area of software architectures and design. He is a member of the IEEE Computer Society.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.

**Luigi Sassoli** received the doctoral degree in informatics engineering and the PhD degree in informatics, multimedia, and telecommunications engineering from the University of Florence in 2002 and 2005, respectively. His research interests mainly focused on state-space analysis methods for performance evaluation and correctness verification of real-time reactive systems, with specific interest on the integration of stochastic characterization into discrete and dense-time models. He was a member of the Software Technologies Lab of the University of Florence from 2002 to 2007. From 2007 to 2008, he worked as a team leader for Gilbarco Veeder-Root. Since 2009, he has been responsible for the POS Support & SQA operative units at Gilbarco Veeder-Root, Florence.

**Laura Carnevali** received the doctoral degree in informatics engineering in 2006 from the University of Florence, where she is now a student in the PhD program in informatics, multimedia, and telecommunications engineering. Her research interests mainly focus on correctness verification and performance evaluation of real-time reactive systems, with specific interest in stochastic characterization of timed models. Her activity also addresses the integration of formal methods in the development life cycle of real-time software. She is a member of the Software Technologies Lab of the University of Florence.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.

Authorized licensed use limited to: Universita degli Studi di Firenze. Downloaded on October 13, 2009 at 02:56 from IEEE Xplore. Restrictions apply.