Partial stochastic characterization of timed runs over DBM domains

L. Carnevali, L. Ridi, E. Vicario
Dipartimento di Sistemi e Informatica - Universitá di Firenze
{carnevali, ridi, vicario}@dsi.unifi.it

***Abstract*— We propose an extension of Time Petri Nets (TPNs) that introduces a partial stochastic characterization of timers, providing a model that regards controllable and non-controllable timers as non-deterministic and stochastic variables, respectively. This induces a partial stochastic characterization of timings in the symbolic state space of the model and supports the evaluation of the probability of execution of any symbolic run as a function of non-deterministic values.**

***Index Terms*—partially stochastic Time Petri Nets, non-Markovian stochastic Petri Nets, symbolic state space analysis, Difference Bound Matrix, stochastic state classes.**

I. INTRODUCTION

The analysis of timed systems can be largely supported by the use of formalisms like Timed Automata or Time Petri Nets (TPNs), whose capabilities of exhaustive state space enumeration and explicit timing parameters manipulation can be exploited in order to validate the correctness of the system and ensure the preservation of concurrency properties. In this context, the identification of the ranges of timing parameters that lead the system to execute along a selected path of the state space represents a problem whose importance is crucial for the observation of behaviors that may exhibit pathological issues. This topic has been developed in [13] in the context of TPNs and directly applied to a real-time testing scenario in [7], by operating a distinction between controllable and non-controllable events, whose timing values are regarded as non-deterministic variables.

The introduction of a stochastic characterization over the space of temporal variables of the model generally highlights two different problems, that in turn require the adoption of different analysis techniques: namely, the evaluation of the probability of a given path and the derivation of a probability distribution over the whole set of timers activated along a path. The former aspect has been dealt with in [10] and [5], with reference to a variant of Timed Automata called Stochastic Timed I/O Automata (STIOA) and to a non-Markovian stochastic Petri Net named stochastic Time Petri Net (sTPN), respectively. In particular, [10] evaluates the probability of a single path, while in [5] the DTMC underlying the model is enumerated, thus providing a wider solution that also encompasses cyclic behaviors. The latter problem is faced within [6] and [12] for sTPN models, and within [14] for STIOA. Both approaches account for the distinction between controllable and non-controllable events in the model; the latter exploits the controllability of certain *input* parameters in order to maximize the probability of execution of the path, under the assumption that all transitions are associated with uniform distributions.

In this paper, we propose partially stochastic Time Petri Nets (psTPNs), which extend TPNs by distinguishing controllable and non-controllable transitions and by associating times to fire of non-controllable transitions with a probability distribution. This induces a partial measure of probability over the space of feasible behaviors of the model, and in particular over the range of feasible timings along a path in the graph of symbolic state classes [3] [13]. We characterize this measure to express the probability distribution of non-controllable times-to-fire as a function of the values taken by the controllable ones. The formulation that we propose here is wider than that of [14] in various aspects: *i)* it considers models where also controllable transitions are bound to fire within constrained intervals; *ii)* it assumes that times to fire of non-controllable transitions are distributed according to any GEN distribution, making explicit effects of the derivation that remain hidden in the case of uniform distributions assumed in [14] and [10]; *iii)* most importantly, it makes evident that quantitative evaluation of the probability distribution of non-controllable events must encompass not only the cases that let the net run along the run but also those that let the behavior deviate from it.

The combination of non-deterministic and stochastic behavior relates to various works in the literature. The general aim of the problem formulation reproduces the concept of [1], where a stochastic fault tree is associated with non-deterministic repair strategies so as to represent the space of possible choices for the selection of an optimal maintenance strategy. More generally, the combination of stochastic and non-deterministic behavior relates to the theory of Markov Decision Processes (MDPs) [2] [8]. However, as opposite to MDPs, in the formulation that we address, timers can be non-Markovian variables with GEN distributions, possibly supported over bounded domains that make the state space non-isomorphic to the graph of reachable markings of the underlying untimed model. Moreover, non-determinism choices do not select discrete (and finite) switches among immediate transitions, but they rather select times to fire taking values within continuous (possibly bounded) domains. As in [4], the joint manipulation of non-deterministic and stochastic parameters unifies qualitative verification of feasible behaviors and quantitative evaluation of their probability. However, while [4] exploits the probabilistic semantics that defines the *likelihood* of each behavior in order to compute the probability of a given property, our effort primarily aims at identifying timed inputs that maximize the probability to cover a previously selected run of the model.

The paper is organized as follows. Sect.II presents psTPNs and discusses symbolic analysis of their state space. The identification of timing boundaries for timers along a path in the state space and their partial stochastic characterization are discussed in Sect.III and Sect.IV, respectively. Finally, conclusions are drawn in Sect.V.

II. Partially stochastic Time Petri Nets

A *partially stochastic Time Petri Net* (psTPN) is a TPN where a subset of transitions is associated with a stochastic characterization of times to fire and where choices among firable transitions with the same time to fire are resolved through a random switch. Formally, a psTPN is a tuple:

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

where:

- the first 9 members comprise the basic model of TPNs [13] (the only difference lies in the partition of the set T of transitions in two subsets T^c and T^{nc} of controllable and non-controllable transitions, respectively);
- $\mathcal{C}$ associates each transition with a weight $\mathcal{C}: T \rightarrow \mathbb{R}^+$;
- $\mathcal{F}$ associates each transition $t \in T^{nc}$ with a static probability distribution $F_t()$ defined over its static firing interval $[EFT(t), LFT(t)]$.

The state of a psTPN is a pair $s = \langle m, \tau \rangle$, where $m : P \rightarrow \mathbb{N}$ is a marking and $\tau : T \rightarrow \mathbb{R}_0^+$ associates each transition with a (dynamic) time to fire. As in TPNs, the state evolves according to two clauses of *firability* and *firing*, which are not reported here for space reasons. As a matter of fact, the overall semantics of the model closely reproduces that of TPNs [13], except for the fact that:

- $\mathcal{C}$ defines a measure of probability for the choice among firable transitions:
$$Prob\{t_0 \text{ is selected}\} = \frac{\mathcal{C}(t_0)}{\sum_{t_i \text{ firable}} \mathcal{C}(t_i)}; \quad (2)$$
- $\mathcal{F}$ defines a measure of probability for the times-to-fire sampled by non-controllable transitions at their enabling time:
$$\tau(t_{nc}) \in [EFT(t_{nc}), LFT(t_{nc})]$$
$$P(\tau(t_{nc}) \leq x) = F_{t_{nc}}(x) \quad (3)$$
$$\forall\ t_{nc} \in T^{nc}.$$

We also assume that static probability distributions are absolutely continuous functions that can be expressed as the integral of a density function:

$$F_t(x) = \int_0^x f_t(y)dy \quad (4)$$

Non-deterministic analysis based on the theory of Difference Bound Matrix (DBM) [9][13] supports the identification of the boundaries of the space of feasible timed behaviors. As already mentioned, *i)* controllable and non-controllable transitions can be concurrently enabled within the same state, and *ii)* probabilities on non-deterministic choices pertain to the selection of timers taking values within continuous domains (non-controllable transitions sample the time to fire according to their respective probability distributions). Both of the choices are resolved at enabling time and the state space is covered through *state classes*, each including a set of states with the same marking, but with different valuations of the vector of times to fire distributed within a DBM domain. A state class is a pair:

$$State\ class = \langle m, D \rangle \quad (5)$$

where m is a marking and D is the DBM domain for times to fire of transitions enabled by m.

The dynamic behavior of the system is covered through a so-called AE reachability relation between state classes [11][13]:

Definition II.1. *Given two state classes $S = \langle m, D \rangle$ and $S' = \langle m', D' \rangle$, we say that S' is a successor of S through t_0, and we write $S \stackrel{t_0}{\Longrightarrow} S'$, iff S' contains all and only the states that are reachable from some state collected in S through some feasible firing of t_0.*

Enumeration of the reachability relation $S^p \stackrel{t}{\Longrightarrow} S^c$ among state classes yields a timed transition system, called *state class graph*:

$$State\ class\ graph = \langle V, E \rangle \quad (6)$$

where vertices in V are state classes and edges in $E \subseteq V \times V$ are labeled with a transition $t \in T$. By transitive closure of Def.II.1, if the psTPN is in a marking m and its timers are distributed over a domain D, then a firing sequence ρ can be executed if and only if ρ is a path originating from class $S^0 = \langle m, D \rangle$ in the state class graph. In this case, the state at the end of the sequence has the marking of the class $S^N = \langle m^N, D^N \rangle$ reached by the path, and its timers are distributed over D^N. According to this, any path in the state class graph represents a qualitative sequence of events that can be executed with a continuous set of timings, and it is referred to as *symbolic run*. A symbolic run is identified by a starting class and a sequence of transitions, and the set of its feasible timings turn out to range within a DBM domain.

III. Inspecting timing boundaries along a symbolic run

We reformulate here the process of trace analysis of [13] by identifying in direct manner a set of constraints which are necessary and sufficient for the execution of the run. We use this result to identify an enlarged domain that includes also inconclusive behaviors.

A. Domain of timings along a symbolic run

We consider a symbolic run ρ that originates from class S^0, visits classes S^1 through S^{N-1} and terminates in S^N (see

Fig.1 for an example). To give identity to transitions enabled in multiple classes, a transition t_i newly enabled in S^n will be denoted with t_i^n in S^n and in all the subsequent classes where it is persistent (i.e. t_i^n identifies a *transition activation* in the dynamic evolution of the model). In Fig.1, t_1^0 and t_1^1 denote the activation of transition t_1 in classes S^0 and S^1, respectively; besides, t_4^1 denotes the activation of transition t_4 that is newly enabled in S^1 and persistent until the firing that enters class S^4.

Each transition activation t_i^n is associated with an *absolute virtual firing time* denoted by τ_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 ρ to the firing that enters S^n. We say that τ_i^n is absolute as it is referred to the start time of the run, and that it is virtual, as t_i^n might not come to fire, either because it is disabled or because the sequence ρ terminates. In Fig.1, this is the case of t_1^1, which is disabled by t_2^1.

To represent the relations existing among transition activations, we introduce the following notation: $\iota(n)$ is the index of the transition that enters class S^n; $\nu(n)$ is the index of the class $S^{\nu(n)}$ in which $t_{\iota(n)}$ is newly enabled; $\gamma(j,n)$ is the index of the transition that disables t_j^n; $\delta(j,n)$ is the index of the class $S^{\delta(j,n)}$ in which $t_{\gamma(j,n)}$ is newly enabled. According to this, $t_{\iota(n)}^{\nu(n)}$ denotes the transition that enables t_i^n, while for any transition t_j^n that does not come to fire, $t_{\gamma(j,n)}^{\delta(j,n)}$ denotes the transition that disables t_j^n or terminates the sequence ρ while t_j^n is still persistent. If t_i^n is newly-enabled in S^0 (i.e. $n=0$), we assume $t_{\iota(0)}^{\nu(0)}$ to be the ground reference t_* of the symbolic run ρ, with timer $\tau_{\iota(0)}^{\nu(0)}=0$. In Fig.1: $t_{\iota(1)}^{\nu(1)}=t_1^0$, $t_{\iota(2)}^{\nu(2)}=t_2^1$, $t_{\iota(3)}^{\nu(3)}=t_5^2$, $t_{\iota(4)}^{\nu(4)}=t_3^1$; t_2^0, t_1^1 and t_4^1 are disabled by t_1^0, t_2^1 and t_3^1, respectively, and thus $t_{\gamma(2,0)}^{\delta(2,0)}=t_1^0$, $t_{\gamma(1,1)}^{\delta(1,1)}=t_2^1$ and $t_{\gamma(4,1)}^{\delta(4,1)}=t_3^1$.

According to the semantics of psTPNs, any vector of virtual firing times representing a feasible timing for ρ satisfies the following three constraints:

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

$$EFT(t_i) \le \tau_i^n - \tau_{\iota(n)}^{\nu(n)} \le LFT(t_i) \qquad \forall\ t_i^n \text{ enabled along } \rho; \tag{7}$$

- 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 ρ:

$$\tau_x^n \ge \tau_{\gamma(x,n)}^{\delta(x,n)} \quad \forall\ t_x^n \text{ enabled but not fired along } \rho; \tag{8}$$

- according to the sequencing of ρ, class S^n is entered not later than S^{n+1}:

$$\tau_{\iota(n+1)}^{\nu(n+1)} \ge \tau_{\iota(n)}^{\nu(n)} \quad \forall\ n \in [0, N-1]. \tag{9}$$

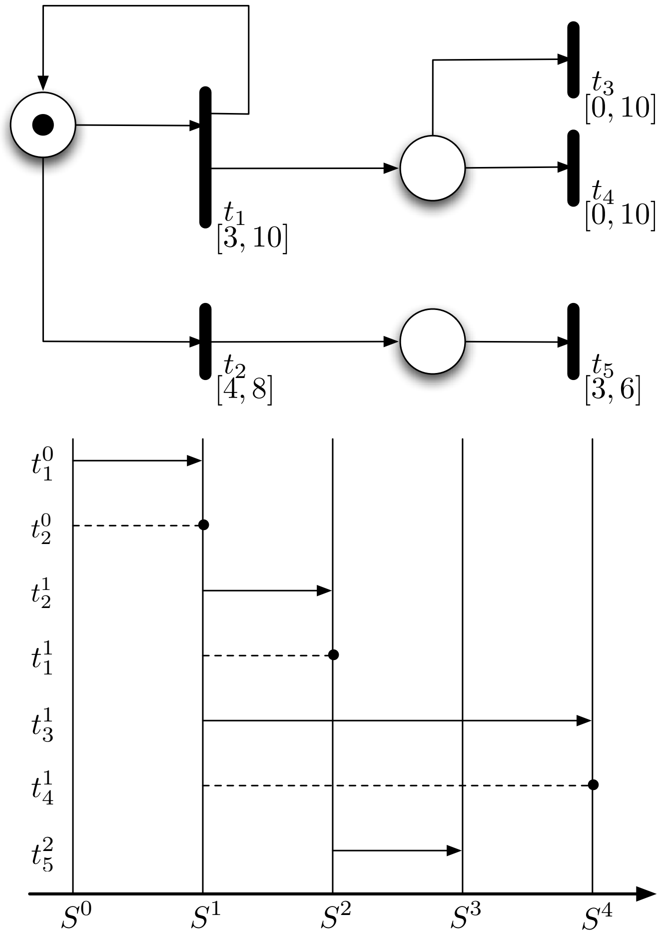

Fig. 1. A simple net and the schema of one of its traces: $\rho = S^0 \xrightarrow{t_1} S^1 \xrightarrow{t_2} S^2 \xrightarrow{t_5} S^3 \xrightarrow{t_3} S^4$. Dotted lines denote times to fire sampled by transitions that do not come to fire; a dot marks the point where the transition is disabled.

Viceversa, it can be proven [6] that constraints in Eqs.(7-9) are also sufficient, provided that all transitions enabled in the starting class S^0 are newly enabled.

According to Eqs.(7-9), the set of all and only the valuations $(\mathbf{x},\mathbf{y})$ of timers $(\boldsymbol{\tau}^c, \boldsymbol{\tau}^{nc})$ that are feasible for ρ is a DBM domain, which we denote by $D_{\boldsymbol{\tau}}$:

$$D_{\boldsymbol{\tau}} = \{(\mathbf{x},\mathbf{y}) \in (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^c}} \times (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^{nc}}} \mid (\mathbf{x},\mathbf{y}) \models (7) \wedge (8) \wedge (9)\} \tag{10}$$

For the purposes of the subsequent treatment, the normal form of $D_{\boldsymbol{\tau}}$ is conveniently represented by distinguishing times to fire of controllable and non-controllable transitions (the normalization of $D_{\boldsymbol{\tau}}$ amounts to an all shortest path problem and can be resolved with complexity $O((N_{\boldsymbol{\tau}^c}+N_{\boldsymbol{\tau}^{nc}})^3)$ [13]). To this end, we denote the set of transition activations t_i^n enabled along ρ as A, and we assume the corresponding virtual firing times τ_i^n be encoded in a vector $\boldsymbol{\tau}$ of length $N_{\boldsymbol{\tau}}$ by means of an indexing function $l(i,n)$ (i.e. $\boldsymbol{\tau}_{l(i,n)} = \tau_i^n$). Moreover, we partition the set of transition activations A into two subsets A^c and A^{nc}, including transition activations with controllable and non-controllable timers, respectively (i.e. $t_i^n \in A^c \Leftrightarrow t_i \in T^c$, $t_i^n \in A^{nc} \Leftrightarrow t_i \in T^{nc}$). In a similar manner, we define a partition of $\boldsymbol{\tau}$ in two sub-vectors $\boldsymbol{\tau}^c$ and $\boldsymbol{\tau}^{nc}$ of length $N_{\boldsymbol{\tau}^c}$ and $N_{\boldsymbol{\tau}^{nc}}$, encoding controllable and non-controllable virtual firing times, respectively (i.e. $\tau_i^n \in \boldsymbol{\tau}^c \Leftrightarrow t_i^n \in A^c$, $\tau_i^n \in \boldsymbol{\tau}^{nc} \Leftrightarrow t_i^n \in A^{nc}$). In addition, a timer τ_i^n is denoted as x_i^n or y_i^n depending on whether it is controllable or non-controllable, respectively (i.e. $\tau_i^n = x_i^n \Leftrightarrow \tau_i^n \in \boldsymbol{\tau}^c$, $\tau_i^n = y_i^n \Leftrightarrow \tau_i^n \in \boldsymbol{\tau}^{nc}$). According to this, the normal form of domain $D_{\boldsymbol{\tau}}$ is expressed as:

$$D_{\boldsymbol{\tau}} = \begin{cases} x_i^n - x_j^m \leq b_{ij}^{nm} \\ x_i^n - y_k^q \leq b_{ik}^{nq} \\ y_h^p - x_j^m \leq b_{hj}^{pm} \\ y_h^p - y_k^q \leq b_{hk}^{pq} \\ \\ \forall\ t_i^n, t_j^m \in A^c,\ t_i^n \neq t_j^m \\ \forall\ t_h^p, t_k^q \in A^{nc},\ t_h^p \neq t_k^q \\ \forall\ m,n,p,q \in [0, N-1] \end{cases} \tag{11}$$

B. Enlarged domain of timings to include inconclusive behaviors

Non-controllable timers are not guaranteed to take values that conform to Eqs.(7-9). Thus, $D_{\boldsymbol{\tau}}$ does not encompass the whole set of timings that may occur during the test run. To this end, we enlarge $D_{\boldsymbol{\tau}}$ in order to derive a domain $D_{\tilde{\boldsymbol{\tau}}}$ that include also those values of non-controllable timers that do not sensitize the test case but may be chosen during execution. According to this, $D_{\tilde{\boldsymbol{\tau}}}$ collects valuations of timers such that values of controllable timers conform to the domain of timings of ρ and values assumed by non-controllable timers only satisfy model constraints.

Formally, $D_{\tilde{\boldsymbol{\tau}}}$ is the intersection between *i)* the projection $D_{\boldsymbol{\tau}^c}$ of domain $D_{\boldsymbol{\tau}}$ on the space of controllable timers $\boldsymbol{\tau}^c$ that identifies all and only the valuations of controllable timers that are feasible for ρ, and *ii)* constraints of Eq.(7) involving at least one non-controllable timer:

$$\begin{aligned} D_{\boldsymbol{\tau}^c} = \quad & \{\mathbf{x} \in (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^c}} \mid \exists\ \mathbf{y} \in (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^{nc}}}\ . \\ & (\mathbf{x}, \mathbf{y}) \models (7) \wedge (8) \wedge (9)\} \end{aligned} \tag{12}$$

$$\begin{aligned} D_{\tilde{\boldsymbol{\tau}}} = \quad & \{(\mathbf{x}, \mathbf{y}) \in (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^c}} \times (\mathbb{R}^+ \cup \{\infty\})^{N_{\boldsymbol{\tau}^{nc}}} \mid \\ & \mathbf{x} \in D_{\boldsymbol{\tau}^c}\ \wedge\ \mathbf{y} \models (7)\} \end{aligned} \tag{13}$$

Fig.2 reports a graphic representation that helps intuition. Based on the properties of the normal representation of a DBM domain [13], the normal form of domain $D_{\boldsymbol{\tau}^c}$ is derived from the normal form of domain $D_{\boldsymbol{\tau}}$ by removing all constraints that involve at least one non-controllable timer:

$$D_{\boldsymbol{\tau}^c} = \begin{cases} x_i^n - x_j^m \leq b_{ij}^{nm} \\ \\ \forall\ t_i^n, t_j^m \in A^c,\ t_i^n \neq t_j^m \\ \forall\ m,n \in [0, N-1] \end{cases} \tag{14}$$

Domain $D_{\tilde{\boldsymbol{\tau}}}$ is obtained by including constraints of Eq.(7) that involve at least one non-controllable timer. We do not report here the explicit representation of $D_{\tilde{\boldsymbol{\tau}}}$; however, since each of its constraints is in DBM form, by construction $D_{\tilde{\boldsymbol{\tau}}}$ also turns out to be a DBM.

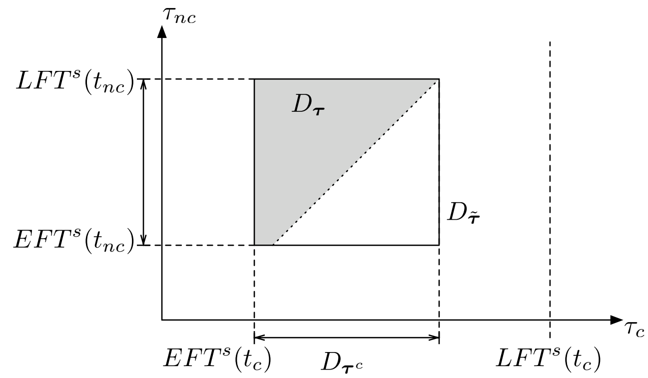

Fig. 2. A bi-variate example with a controllable timer τ_c and a non-controllable timer τ_{nc} (i.e. $\boldsymbol{\tau}^c = (\tau_c)$ and $\boldsymbol{\tau}^{nc} = (\tau_{nc})$). Timers τ_c and τ_{nc} are associated with transitions t_c and t_{nc}, respectively, newly-enabled in S^0. Domain $D_{\boldsymbol{\tau}}$ is the set of feasible timers valuations for $\rho = S^0 \xrightarrow{t_c} S^1 \xrightarrow{t_{nc}} S^2$; $D_{\tilde{\boldsymbol{\tau}}}$ is the corresponding enlarged domain.

IV. Partial stochastic characterization of timings along a symbolic run

$\mathcal{F}$ induces a partial stochastic characterization of timings along a symbolic run ρ and supports the derivation of a measure of probability to actually execute the path as a function of non-deterministic values assumed by controllable timers. To this end, we first identify the domain of non-controllable timers for a given valuation of controllable timers and we then derive the stochastic characterization that this choice induces. For an assigned valuation $\mathbf{x}$ of controllable timers, let $D_{\tilde{\boldsymbol{\tau}}^{nc}}(\mathbf{x})$ and $D_{\boldsymbol{\tau}^{nc}}(\mathbf{x})$ be the projections of domains $D_{\tilde{\boldsymbol{\tau}}}$ and $D_{\boldsymbol{\tau}}$, respectively, over the space of non-controllable timers (see Fig.3). We define a family $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x})$ of functions of non-controllable timers, indexed by valuations $\mathbf{x}$ of controllable timers and supported over domain $D_{\tilde{\boldsymbol{\tau}}^{nc}}(\mathbf{x})$: $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x})$ associates each determination $\mathbf{x}$ of controllable timers with the probability density function of non-controllable timers, providing a measure of the stochastic characterization of non-controllable timers induced by the choice $\mathbf{x}$ of controllable timers (see Fig.3). Non-controllable timers are statistically independent variables that sample their times to fire at newly enabling according to their respective static probability distributions. According to this, the family of functions $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x})$ can be expressed as follows:

$$\begin{aligned} f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x}) \quad &= \prod_{t_i^n \in A^{nc}} f_{t_i}(y_i^n - \tau_{\iota(n)}^{\nu(n)}) \\ &= \prod_{\substack{t_i^n\ \in\ A^{nc} \\ t_{\iota(n)}^{\nu(n)}\ \in\ A^{nc}}} f_{t_i}(y_i^n - y_{\iota(n)}^{\nu(n)}) \\ &\quad \cdot \prod_{\substack{t_i^n\ \in\ A^{nc} \\ t_{\iota(n)}^{\nu(n)}\ \in\ A^c \cup \{t_*\}}} f_{t_i}(y_i^n - x_{\iota(n)}^{\nu(n)}) \end{aligned} \tag{15}$$

For a given valuation $\mathbf{x} = \mathbf{x}_1$ of controllable timers, the integral of function $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x}_1)$ over domain $D_{\boldsymbol{\tau}^{nc}}(\mathbf{x}_1)$ represents the probability to execute the symbolic run ρ under the assumption of the choice $\mathbf{x}_1$ on controllable timers. According to this, the integral of the family of functions $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x})$ over domain $D_{\boldsymbol{\tau}^{nc}}(\mathbf{x})$ defines a new function $f(\mathbf{x})$, that associates each valuation $\mathbf{x}$ of controllable timers with the probability to actually execute the test-case ρ:

$$
\begin{aligned}
f(\mathbf{x}) &= \int_{D_{\boldsymbol{\tau}^{nc}}(\mathbf{x})} f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x}) \quad d(\boldsymbol{\tau}^{nc}) \\
&= Prob\{(\mathbf{x},\mathbf{y}) \in D_{\boldsymbol{\tau}} \mid \boldsymbol{\tau}^c = \mathbf{x}\} \\
&= Prob\{\rho \text{ is executed} \mid \boldsymbol{\tau}^c = \mathbf{x}\}
\end{aligned} \tag{16}
$$

Function $f(\mathbf{x})$ provides the probability of successful sensitization as a function of non-deterministic values assumed by controllable timers, in the light of the stochastic characterization of non-controllable timers.

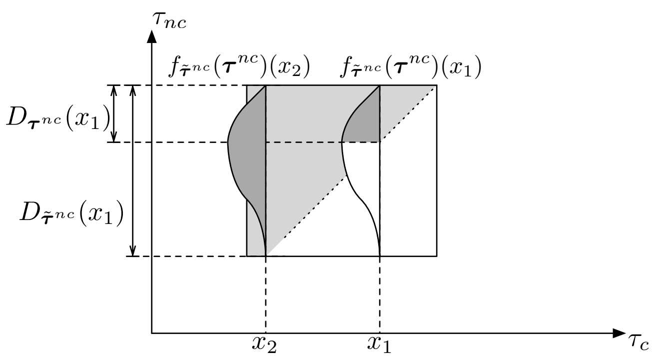

Fig. 3. Referring to the example of Fig.2, the picture highlights *i)* projections $D_{\boldsymbol{\tau}^{nc}}(x_1)$ and $D_{\tilde{\boldsymbol{\tau}}^{nc}}(x_1)$ of domains $D_{\boldsymbol{\tau}}$ and $D_{\tilde{\boldsymbol{\tau}}}$, respectively, over the space of non-controllable timer τ_{nc} for valuation x_1 of controllable timer τ_c; *ii)* two instances $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(x_1)$ and $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(x_2)$ of the family of functions $f_{\tilde{\boldsymbol{\tau}}^{nc}}(\boldsymbol{\tau}^{nc})(\mathbf{x})$ for valuations $\mathbf{x} = (x_1)$ and $\mathbf{x} = (x_2)$ of the controllable timer τ_c, respectively.

V. Conclusions

We considered an extension of TPNs where a subset of transitions are associated with a stochastic characterization of the time to fire sampled at the enabling time. In a pragmatic perspective, transitions associated with this stochastic characterization represent non-controllable temporal parameters, while transitions that are left non-deterministic account for parameters that can be controlled in the design stage. This induces a partial measure of probability over the space of feasible timings that can occur in the execution of a symbolic run and supports the identification of a measure of probability of successful sensitization as a function of non-deterministic values assumed by controllable timers.

The result opens the way to further theoretical problems and practical applications. In the theoretical perspective, the identification of the probability density function $f(\mathbf{x})$ opens the way to a problem of optimization, aimed at the identification of valuations of controllable timers that maximize the probability density function $f(\mathbf{x})$ of successful sensitization. In the practical perspective, the approach finds a direct application in real-time testing, where $f(\mathbf{x})$ can be employed in the selection of timed inputs through a method for the simulation of multivariate distributions.

References

[1] M. Beccuti, D. Codetta-Raiteri, G. Franceschinis, and S. Haddad. Non deterministic repairable fault trees for computing optimal repair strategy. In *Proc. of the* 3^{rd} *Int. Conf. on Performance Evaluation Methodologies and Tools (VALUETOOLS'08)*, Athens, Greece, October 2008.

[2] M. Beccuti, G. Franceschinis, and S. Haddad. Markov decision petri net and markov decision well-formed net formalisms. In *Petri Nets and Other Models of Concurrency ICATPN 2007*, 2007.

[3] B. Berthomieu and M. Diaz. Modeling and verification of time dependent systems using time petri nets. *IEEE Trans. on Soft. Eng.*, 17(3), March 1991.

[4] N. Bertrand, P. Bouyer, T. Brihaye, and N. Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *Proc. of the* 5^{th} *Int. Conf. on Quantitative Evaluation of Systems (QEST '08)*, pages 55–64, Sept. 2008.

[5] G. Bucci, R.Piovosi, L. Sassoli, and E. Vicario. Introducing probability within state class analysis of dense time dependent systems. *Proc. of the* 2^{nd} *Int. Conf. on Quant. Evaluation of Sys. (QEST '05)*, Sept. 2005.

[6] L. Carnevali, L. Sassoli, and E. Vicario. Using stochastic state classes in quantitative evaluation of dense-time reactive systems. accepted for publication in the IEEE Trans. on Soft. Eng.

[7] L. Carnevali, L. Sassoli, and E. Vicario. Sensitization of symbolic runs in real-time testing using the oris tool. *Proc. of the IEEE Conf. on Emerging Technologies and Factory Automation (ETFA)*, Sept. 2007.

[8] L. de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In *CONCUR '99: Proceedings of the* 10^{th} *International Conference on Concurrency Theory*, pages 66–81, London, UK, 1999. Springer–Verlag.

[9] D. Dill. Timing assumptions and verification of finite-state concurrent systems. *Proc. Workshop on Computer Aided Verification Methods for Finite State Systems*, 1989.

[10] M. Jurdziński, D. Peled, and H. Qu. Calculating probabilities of real-time test cases. In *Proceedings of the 5th International Workshop on Formal approaches to software testing (FATES '05)*, pages 134–151, 2005.

[11] W. Penczek and A. Polrola. Specification and model checking of temporal properties in time petri nets and timed automata. *Proocedings of the* 25^{th} *Int. Conf on Application and Theory of Petri Nets, ICATPN2004*, June 2004.

[12] L. Sassoli and E. Vicario. Close form derivation of state-density functions over dbm domains in the analysis of non-markovian models. *Proc. of the* 4^{nd} *Int. Conf. on the Quant. Evaluation of Sys. (QEST)*, September 2007.

[13] E. Vicario. Static analysis and dynamic steering of time dependent systems using time petri nets. *IEEE Trans. on Soft. Eng.*, August 2001.

[14] N. Wolovick, P.R. D'Argenio, and H. Qu. Optimizing probabilities of real-time test case execution. In *Software Testing Verification and Validation, 2009. ICST '09. International Conference on*, pages 1–10, April 2009.