Comparison of Different Semantics for Time Petri Nets

Béatrice Berard, Franck Cassez, Serge Haddad, Didier Lime, Olivier Henri Roux

► To cite this version:

Béatrice Berard, Franck Cassez, Serge Haddad, Didier Lime, Olivier Henri Roux. Comparison of Different Semantics for Time Petri Nets. 3rd Int. Symp. on Automated Technology for Verification and Analysis (ATVA'05), Oct 2005, Taiwan, Taiwan. pp.293–307. inria-00368580

HAL Id: inria-00368580

<https://hal.inria.fr/inria-00368580>

Submitted on 17 Mar 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparison of Different Semantics for Time Petri Nets

B. Bérard¹, F. Cassez², S. Haddad¹, Didier Lime³, O.H. Roux²

¹ LAMSADE, Paris, France

E-mail: {beatrice.berard | serge.haddad}@lamsade.dauphine.fr

² IRCCyN, Nantes, France

{Franck.Cassez | Olivier-h.Roux}@irccyn.ec-nantes.fr

³ CISS, Aalborg, Denmark

Didier@cs.aau.dk

Abstract. In this paper we study the model of Time Petri Nets (TPNs) where a time interval is associated with the firing of a transition, but we extend it by considering general intervals rather than closed ones. A key feature of timed models is the memory policy, i.e. which timing informations are kept when a transition is fired. The original model selects an *intermediate* semantics where the transitions disabled after consuming the tokens, as well as the firing transition, are reinitialised. However this semantics is not appropriate for some applications. So we consider here two alternative semantics: the *atomic* and the *persistent atomic* ones. First we present relevant patterns of discrete event systems which show the interest of these semantics. Then we compare the expressiveness of the three semantics w.r.t. the weak time bisimilarity establishing inclusion results in the general case. Furthermore we show that some inclusions are strict with unrestricted intervals even when nets are bounded. Then we focus on bounded TPNs with upper-closed intervals and we prove that the semantics are equivalent. Finally taking into account both the practical and the theoretical issues, we conclude that persistent atomic semantics should be preferred.

Keywords: Time Petri Nets, Timed Bisimilarity, Expressiveness.

1 Introduction

Since their introduction, Petri nets have been successfully applied for the design and analysis of discrete event systems. However with the development of critical systems, time has become a significant issue for their correctness. So numerous timed extensions have been proposed for Petri nets. These extensions can be roughly divided into three categories w.r.t. their application field:

- *Timed Petri nets* [17] include a duration associated with each transition in order to model for instance scheduling policies in production management;
- *Stochastic Petri nets* [2] include a probability distribution associated with each transition in order to evaluate the transient or steady-state behaviour of a system where duration of actions are obtained by statistical observations;

- *Time Petri nets* [15] include a time interval associated with each transition (specifying the possible time elapsing before firing) in order to model systems which may evolve in a non deterministic way.

The latter model has a lot of alternatives: for instance time constraints may be associated to places [12] or to arcs [1,9], timed synchronisations may be added [10], etc. Here we stick to the original model for which significant theoretical developments [5,16,13,7] have been obtained leading to efficient verification tools on models of large systems [6,11]. Moreover when bounded this model can be translated into timed automata [8] which are also extensively studied with successful applications [3].

More precisely we consider a slight extension of TPNs where any kind of time interval can be associated with the firing of a transition instead of the closed ones in the original definition. Our work focuses on the impact of the memory policy on the expressiveness of the model. The memory policy specifies which timing informations are kept when a transition is fired. The original model selects an intermediate semantics meaning that the transitions disabled after consuming the tokens, and the firing transition, are reinitialised. Here we propose two alternative semantics: the atomic and the persistent atomic ones. We first present significant examples where these semantics are more appropriate than the original one. Then we compare the expressiveness of the three semantics w.r.t. the weak time bisimilarity which is the standard equivalence relation used for such comparisons. First we establish inclusion results in the general case. Furthermore we show that some inclusions are strict with unrestricted intervals even when nets are bounded. The key point of the previous result is the presence of non *upper-closed* intervals. Then we focus on bounded TPNs with upper-closed intervals and we establish that the semantics are equivalent. All our translations are optimal since the size of the produced net is linear w.r.t. the size of the emulated net.

The paper is organised as follows. Section 2 introduces the syntax and semantics of TPNs, timed transition systems and timed bisimilarity. Section 3 is devoted to the comparison between the different semantics for TPNs with general intervals both in the bounded and the unbounded case. In section 4, we prove that the three semantics are equivalent for bounded TPNs with upper-closed intervals. Finally we conclude in section 5.

2 Time Petri Nets

Notations. Let Σ be a finite set (or alphabet). Σ^* denotes the set of finite words on Σ . If $w = a_1 \cdots a_n$, the *length* of w denoted $|w|$, is n . We also use $\Sigma_\varepsilon = \Sigma \cup \{\varepsilon\}$ with $\varepsilon \notin \Sigma$, where ε is the empty word. B^A stands for the set of mappings from A to B . If A is finite and $|A| = n$, an element of B^A is also a vector in B^n . The usual operators $+$, $-$, $<$ and $=$ are used on vectors of A^n with $A = \mathbb{N}, \mathbb{Q}, \mathbb{R}$ and are the point-wise extensions of their counterparts in A . The set $\mathbb{B}$ denotes the boolean values $\{\text{tt}, \text{ff}\}$ and $\mathbb{R}_{\geq 0}$ denotes the set of non negative reals. A *valuation* ν over a set of variables X is an element of $\mathbb{R}_{\geq 0}^X$. For $\nu \in \mathbb{R}_{\geq 0}^X$

and $d \in \mathbb{R}_{\geq 0}$, $\nu + d$ denotes the valuation defined by $(\nu + d)(x) = \nu(x) + d$. $\mathbf{0}$ denotes the valuation s.t. $\forall x \in X, \nu(x) = 0$.

An interval I is a $\mathbb{Q}_{\geq 0}$ -interval of $\mathbb{R}_{\geq 0}$ iff its left endpoint belongs to $\mathbb{Q}_{\geq 0}$ and its right endpoint belongs to $\mathbb{Q}_{\geq 0} \cup \{\infty\}$. We let $I^\downarrow = \{x \mid \exists y \in I \wedge y \geq x\}$ be the *downward closure* of I . We denote by $\mathcal{I}(\mathbb{Q}_{\geq 0})$ the set of $\mathbb{Q}_{\geq 0}$ -intervals of $\mathbb{R}_{\geq 0}$.

2.1 Timed Transition Systems (TTS) and weak time bisimulation

Timed transition systems describe systems which combine discrete and continuous evolutions. We present here a standard version.

Definition 1 (Timed Transition Systems). A timed transition system (TTS) over the set of actions Σ_ε is a tuple $S = (Q, q_0, \Sigma_\varepsilon, \longrightarrow)$ where:

- Q is a set of states,
- $q_0 \in Q$ is the initial state,
- Σ_ε is a finite set of actions disjoint from $\mathbb{R}_{\geq 0}$,
- $\longrightarrow \subseteq Q \times (\Sigma_\varepsilon \cup \mathbb{R}_{\geq 0}) \times Q$ is a set of edges. If $(q, e, q') \in \longrightarrow$, we also write $q \xrightarrow{e} q'$. For a transition $q \xrightarrow{d} q'$ with $d \in \mathbb{R}_{\geq 0}$, the value d represents a relative time stamp.

We make the following common assumptions about TTS:

- TIME-DETERMINISM: if $q \xrightarrow{d} q'$ and $q \xrightarrow{d} q''$ with $d \in \mathbb{R}_{\geq 0}$, then $q' = q''$,
- 0-DELAY: $q \xrightarrow{0} q$,
- ADDITIVITY: if $q \xrightarrow{d} q'$ and $q' \xrightarrow{d'} q''$ with $d, d' \in \mathbb{R}_{\geq 0}$, then $q \xrightarrow{d+d'} q''$,
- CONTINUITY: if $q \xrightarrow{d} q'$, then for every d' and d'' in $\mathbb{R}_{\geq 0}$ such that $d = d' + d''$, there exists q'' such that $q \xrightarrow{d'} q'' \xrightarrow{d''} q'$.

A run ρ of length $n \geq 0$ is a finite sequence of transitions of the form

$$\rho = q_0 \xrightarrow{d_0} q'_0 \xrightarrow{a_0} q_1 \xrightarrow{d_1} q'_1 \xrightarrow{a_1} \dots q_n \xrightarrow{d_n} q'_n$$

where discrete actions alternate with (possibly null) durations. We also write this run as $q \xrightarrow{d_0 a_0 \dots d_n} q'$. $\text{Untimed}(\rho)$ is the word of Σ^* obtained by concatenation of labels $a_0, \dots, a_{n-1}$ (remember that ε is the empty word). $\text{Duration}(\rho) = \sum_{i=0}^n d_i$. Furthermore, a run can be safely by inserting or deleting a null duration step and splitting a duration step or merging two consecutive duration steps as long as the total duration is the same.

It is well-known that the concept of weak timed bisimilarity is central among equivalence relations between systems since, for instance, two TTS are not distinguishable by formulas of most common timed arborescent temporal logics.

Let $S = (Q, q_0, \Sigma_\varepsilon, \longrightarrow)$ be a TTS. We define the relation $\longrightarrow \triangleright \subseteq Q \times (\Sigma \cup \mathbb{R}_{\geq 0}) \times Q$ by:

- $q \xrightarrow{d} q'$ iff there is a run $\rho = q \xrightarrow{\sigma} q'$ with $\text{Untimed}(\rho) = \varepsilon$ and $\text{Duration}(\rho) = d$,
- $q \xrightarrow{a} q'$ with $a \in \Sigma$ iff there is a run $\rho = q \xrightarrow{\sigma} q'$ with $\text{Untimed}(\rho) = a$ and $\text{Duration}(\rho) = 0$,

Definition 2 (Weak Timed Bisimilarity). Let $S_1 = (Q_1, q_0^1, \Sigma_\varepsilon, \longrightarrow_1)$ and $S_2 = (Q_2, q_0^2, \Sigma_\varepsilon, \longrightarrow_2)$ be two TTS and $\approx$ be a binary relation over $Q_1 \times Q_2$. We write $q \approx q'$ for $(q, q') \in \approx$. $\approx$ is a weak timed bisimulation relation between S_1 and S_2 if:

- $q_0^1 \approx q_0^2$;
- if $q_1 \xrightarrow{a} q'_1$ with $a \in \Sigma \cup \mathbb{R}_{\geq 0}$ and $q_1 \approx q_2$ then $\exists q'_2 \xrightarrow{a} q'_2$ such that $q'_1 \approx q'_2$; conversely if $q_2 \xrightarrow{a} q'_2$ with $a \in \Sigma \cup \mathbb{R}_{\geq 0}$ and $q_1 \approx q_2$ then $\exists q'_1 \xrightarrow{a} q'_1$ such that $q'_1 \approx q'_2$.

Two TTS S_1 and S_2 are weakly timed bisimilar if there exists a weak timed bisimulation relation between S_1 and S_2 . We write $S_1 \approx_{\mathcal{W}} S_2$ in this case.

Let $\mathcal{C}$ and $\mathcal{C}'$ be two classes of TTS. The next definition formalises the relative expressiveness of $\mathcal{C}$ and $\mathcal{C}'$.

Definition 3 (Expressiveness w.r.t. Weak Timed Bisimilarity). The class $\mathcal{C}$ is more expressive than $\mathcal{C}'$ w.r.t. weak timed bisimilarity if for all $B' \in \mathcal{C}'$ there is a $B \in \mathcal{C}$ s.t. $B \approx_{\mathcal{W}} B'$. We write $\mathcal{C}' \leq_{\mathcal{W}} \mathcal{C}$ in this case. If moreover there is a $B \in \mathcal{C}$ s.t. there is no $B' \in \mathcal{C}'$ with $B \approx_{\mathcal{W}} B'$, then $\mathcal{C}' <_{\mathcal{W}} \mathcal{C}$. If both $\mathcal{C}' \leq_{\mathcal{W}} \mathcal{C}$ and $\mathcal{C} \leq_{\mathcal{W}} \mathcal{C}'$ then $\mathcal{C}$ and $\mathcal{C}'$ are equally expressive w.r.t. weak timed bisimilarity, and we write $\mathcal{C} \approx_{\mathcal{W}} \mathcal{C}'$.

2.2 Time Petri Nets

Time Petri Nets (TPN) were introduced in [15] and extend Petri Nets with timing constraints on the firings of transitions. In a TPN, a time interval is associated with each transition. An implicit clock can then be associated with each enabled transition, and gives the elapsed time since it was last enabled. An enabled transition can be fired if its clock value belongs to the interval of the transition. Furthermore, time cannot progress if time elapsing would result in leaving the interval of a transition. The following definitions formalise these principles.

Definition 4 (Labeled Time Petri Net). A Labeled Time Petri Net $\mathcal{N}$ is a tuple $(P, T, \Sigma_\varepsilon, \bullet(\cdot), (\cdot)^\bullet, M_0, A, I)$ where:

- P is a finite set of places;
- T is a finite set of transitions with $P \cap T = \emptyset$;
- $\bullet(\cdot) \in (\mathbb{N}^P)^T$ is the backward incidence mapping; $(\cdot)^\bullet \in (\mathbb{N}^P)^T$ is the forward incidence mapping;
- $M_0 \in \mathbb{N}^P$ is the initial marking;

- $\Lambda : T \rightarrow \Sigma_\varepsilon$ is the labeling function;
- $I : T \rightarrow \mathcal{I}(\mathbb{Q}_{\geq 0})$ associates with each transition a firing interval;

We also use $\bullet t$ (resp. $t^\bullet$) to denote the set of places $\bullet t = \{p \in P \mid \bullet t(p) > 0\}$ (resp. $t^\bullet = \{p \in P \mid t^\bullet(p) > 0\}$) as it is common in the literature⁴.

The net of figure 1 illustrates the graphical representation of a TPN. Each transition is represented with its label and its interval. For instance transition t_1 has label λ_1 and interval $[1, \infty[$.

Fig. 1. An example of TPN

Semantics of Time Petri Nets. The semantics of TPNs is given in terms of Timed Transition Systems. A *marking* M of a TPN is a mapping in $\mathbb{N}^P$ and $M(p)$ is the number of tokens in place p . A transition t is *enabled* in a marking M iff $M \geq \bullet t$. We denote by $En(M)$ the set of enabled transitions in M . To decide whether a transition t can be fired, we need to know for how long it has been enabled: if this amount of time lies within the interval $I(t)$, t can actually be fired, otherwise it cannot. On the other hand time can progress only if the enabling duration still belongs to the downward closure of the interval associated with an enabled transition. Let $\nu \in (\mathbb{R}_{\geq 0})^{En(M)}$ be a *valuation* such that each value $\nu(t)$ is the time elapsed since transition t was last enabled. A *configuration* of the TPN $\mathcal{N}$ is a pair (M, ν) . An *admissible configuration* of a TPN is a configuration (M, ν) s.t. $\forall t \in En(M), \nu(t) \in I(t)^\downarrow$. We let $ADM(\mathcal{N})$ be the set of admissible markings.

When defining the semantics of a TPN, three kinds of policies must be fixed.

The choice policy concerns the choice of the next event to be fired (scheduled). For TPNs (and also timed automata), this choice is non deterministic (possible alternatives use priorities, probabilities, etc.).

The service policy concerns the possibility of simultaneous instances of a same event to occur. In the context of Petri nets, this is formalised by the enabling *degree* of a transition. Here we adopt the *single-server policy* (at most one instance of a firing per transition in every state). Our results could be extended to the multiple server at the price of intricate notations.

⁴ Whether $\bullet t$ (resp. $t^\bullet$) stands for a vector of $(\mathbb{N}^P)^T$ or a subset of P will be unambiguously defined by the context.

The memory policy concerns the updating of timing informations when a discrete step occurs. The key issue in the semantics is to define when to reset the clock measuring the time since a transition was last enabled. This can only happen when a transition is fired. We let $\uparrow enabled(t', M, t) \in \mathbb{B}$ be true if t' is *newly enabled* by the firing of transition t from marking M , and false otherwise.

Let M be a marking and $t \in En(M)$. The firing of t leads to a new marking $M' = M - \bullet t + t^\bullet$. Three semantics are possible:

I: The *intermediate semantics* (*I*) considers that the firing of a transition is performed in two steps: consuming the input tokens in $\bullet t$, and then producing output tokens in $t^\bullet$. The fact that a transition t' is newly enabled on the firing of a transition $t \neq t'$ is determined w.r.t. the intermediate marking $M - \bullet t$. When a transition t is fired it is newly enabled whatever the intermediate marking. We denote by $\uparrow enabled_I(t', M, t)$ the newly enabled predicate in this case. This mapping is defined by:

$$\begin{aligned} \uparrow enabled_I(t', M, t) = & (t' \in En(M - \bullet t + t^\bullet)) \\ & \wedge (t' \notin En(M - \bullet t) \vee (t = t')) \end{aligned} \quad (1)$$

A: The *atomic semantics* considers that the firing of a transition is obtained by an atomic step. The corresponding mapping $\uparrow enabled_A(t', M, t)$ is defined by:

$$\uparrow enabled_A(t', M, t) = (t' \in En(M - \bullet t + t^\bullet)) \wedge (t' \notin En(M) \vee (t = t')) \quad (2)$$

PA: The *persistent atomic semantics* considers that the firing of a transition is also obtained by an atomic step. The difference with the *A* semantics is only on the value of $\uparrow enabled_A(t', M, t)$ when $t = t'$. The fired transition here is handled as any other one:

$$\uparrow enabled_{PA}(t', M, t) = t' \in En(M - \bullet t + t^\bullet) \wedge (t' \notin En(M)) \quad (3)$$

Note that we have the relation: $(3) \Rightarrow (2) \Rightarrow (1)$ but as we shall see on the example this does not imply any inclusion relation between the different behaviours.

We now define the semantics of a TPN, which is parameterised by the choice of the $\uparrow enabled$ predicate.

Definition 5 (Semantics of TPN). Let $s \in \{I, A, PA\}$. The s -semantics of a TPN $\mathcal{N} = (P, T, \Sigma_\varepsilon, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda, I)$ is a timed transition system $S_{\mathcal{N}} = (Q, q_0, T, \rightarrow)$ where: $Q = ADM(\mathcal{N})$, $q_0 = (M_0, \mathbf{0})$, and $\rightarrow \in Q \times (\Sigma_\varepsilon \cup \mathbb{R}_{\geq 0}) \times Q$ consists of the discrete and continuous transition relations:

- $\forall (M, \nu) \in ADM(\mathcal{N}), \forall t \in En(M), s.t. \nu(t) \in I(t)$
the discrete transition relation is defined by:

$$(M, \nu) \xrightarrow{\Lambda(t)} (M - \bullet t + t^\bullet, \nu') \text{ where } \forall t \in En(M - \bullet t + t^\bullet),$$

$$\nu'(t) = \begin{cases} 0 & \text{if } \uparrow enabled_s(t', M, t), \\ \nu(t) & \text{otherwise.} \end{cases}$$

- $\forall (M, \nu) \in ADM(\mathcal{N}), \forall d \in \mathbb{R}_{\geq 0}, s.t. \forall t \in En(M), \nu(t) + d \in I(t)^\downarrow$, the continuous transition relation is defined by: $(M, \nu) \xrightarrow{d} (M, \nu + d)$

We simply write $(M, \nu) \xrightarrow{w}$ to emphasise that a sequence of transitions w can be fired in $S_{\mathcal{N}}$ from (M, ν) . If $Duration(w) = 0$ we say that w is an *instantaneous firing sequence*. The set of *reachable markings* of $\mathcal{N}$ is $Reach(\mathcal{N}) = \{M \in \mathbb{N}^P \mid \exists (M, \nu) \mid (M_0, \mathbf{0}) \xrightarrow{w} (M, \nu)\}$. A net is bounded iff there exists an integer B such that $\forall M \in Reach(\mathcal{N}), \forall p \in P, M(p) \leq B$.

We illustrate the different semantics on the example of fig. 1. $(M_0, \mathbf{0}) \xrightarrow{1t_1t_1}$ is a run for the *PA* semantics but not for the other ones since the second firing of t_1 should imply a delay of at least 1 time unit. $(M_0, \mathbf{0}) \xrightarrow{1t_1t_2}$ is a run for the *PA* and *A* semantics but not for the *I* semantics since as t_1 consumes (and produces again) the token in place p_2 this should imply a delay of at least 1 time unit before the firing of t_2 . $(M_0, \mathbf{0}) \xrightarrow{1t_11t_2}$ is a run for the *I* semantics but not for the other ones since after t_1 fires, t_2 is not newly enabled and time cannot elapse before it is either fired or disabled.

The intermediate semantics *I*, based on [4,5] is the most common one. However, we provide two significant patterns (among other ones) of discrete event systems where the other semantics are more appropriate. Consider the net of figure 2 which models a component whose state is checked by an observer in order to react (by firing transition t). Let us emphasise that this modeling by a loop is standard in the untimed context. This observation does not interfere with the behaviour of the component using *A* and *PA* semantics while the *I* semantics reinitialises the delay of transitions t_1 and t_2 .

Fig. 2. A case where *PA* or *A* are more convenient than *I*: observation of a component

The subnet of fig. 3 models clients waiting for some information to be produced by a source (in 1 time unit) and then sent to every client (in a negligible time). The two firings of t are performed at the same time only with the *PA* semantics.

The comparison of expressive powers of the three semantics has not been investigated and is the topic of our work.

Fig. 3. A case where PA is more convenient than A and I : instantaneous multicast

3 Inclusion results for the semantics I , A and PA

3.1 Large inclusion results

We now establish inclusion relations between the three semantics for TPNs.

In order to alleviate the figures, transitions are filled in black when their firing interval is $[0, 0]$ and we omit their label when equal to ε .

Proposition 1. *Let $\mathcal{N}$ be a time Petri net with intermediate semantics. There exists a TPN $\overline{\mathcal{N}}$ with atomic semantics which is weakly timed bisimilar to $\mathcal{N}$. The size of $\overline{\mathcal{N}}$ is linear w.r.t. the size of $\mathcal{N}$. Furthermore if $\mathcal{N}$ is bounded then $\overline{\mathcal{N}}$ is bounded.*

Proof. The construction is quite easy. Let $T' \subseteq T$ be the subset of transitions which have at least one input place. The set of places of $\overline{\mathcal{N}}$ is obtained by adding to the set of places of $\mathcal{N}$ a new place for each transition t from $\mathcal{N}$: $\overline{P} = P \cup \{p_t \mid t \in T'\}$. The transitions of $T \setminus T'$ are unchanged. The transitions T' of $\mathcal{N}$ are duplicated in $\overline{\mathcal{N}}$: $\overline{T} = T'^+ \cup T'^- \cup (T \setminus T')$ and the construction follows Figure 4, from left to right.

Fig. 4. From I to A

We consider the equivalence relation $\mathcal{R}$ which contains all pairs $((M, \nu), (\overline{M}, \overline{\nu}))$ such that:

- for all $p \in P$, $M(p) = \overline{M}(p) + \sum_{t \in T'} t^\bullet(p) \cdot \overline{M}(p_t)$
- for all $t \in T \setminus T'$, $\nu(t) = \overline{\nu}(t)$ and for all $t \in T' \cap \text{En}(M)$, $\nu(t) = \overline{\nu}(t^-)$ if t^- is enabled in $\overline{M}$ and 0 otherwise. The latter case corresponds in $\mathcal{N}$ to a newly enabled transition.

To prove that $\mathcal{R}$ is a bisimulation, we first note that, with the definition above for markings, from any configuration $(\overline{M}, \overline{\nu})$, we can reach instantaneously a configuration $(\overline{M}_1, \overline{\nu}_1)$ such that $\overline{M}_1(p_t) = 0$ for all $t \in T'$, with the firing of a (possibly empty) sequence of transitions in T^+ . Moreover, the relation between valuations implies that $(\overline{M}_1, \overline{\nu}_1)$ is still equivalent to (M, ν) .

Consider now a pair $((M, \nu), (\overline{M}, \overline{\nu})) \in \mathcal{R}$.

- if $(M, \nu) \xrightarrow{t} (M', \nu')$ with $t \in T'$, then from the remark above, we first fire a sequence from $\overline{M}$ to empty all places $p_{t'}$, leading instantaneously to $(\overline{M}_1, \overline{\nu}_1)$, which is equivalent to (M, ν) . Then transition t^- can be fired from $(\overline{M}_1, \overline{\nu}_1)$, immediately followed by t^+ , leading to $(\overline{M}', \overline{\nu}')$, where all places p_t are empty again. Moreover, the transitions newly enabled by t^+ in $\overline{\mathcal{N}}$ are exactly those which were newly enabled by t in $\mathcal{N}$, so that $(M', \nu') \mathcal{R} (\overline{M}', \overline{\nu}')$.
- Conversely, suppose that a transition is fired from $(\overline{M}, \overline{\nu})$ in $\overline{\mathcal{N}}$. If the transition is some t^+ , then the new configuration $(\overline{M}_1, \overline{\nu}_1)$ is still equivalent to (M, ν) (as above), thus no move at all is necessary in $\mathcal{N}$.

If $(\overline{M}, \overline{\nu}) \xrightarrow{t^-} (\overline{M}', \overline{\nu}')$, then t can be fired from (M, ν) and the resulting marking is, (M', ν') , equivalent to $(\overline{M}', \overline{\nu}')$.

- if $(M, \nu) \xrightarrow{d} (M, \nu + d)$, for some delay d , then again we have to apply the emptying sequence from $(\overline{M}, \overline{\nu})$, to reach a configuration $(\overline{M}_1, \overline{\nu}_1)$ still equivalent to (M, ν) , where time can elapse. The relation between ν and $\overline{\nu}_1$ implies that this is possible, leading to $(\overline{M}_1, \overline{\nu}_1 + d)$.
- Conversely, if $(\overline{M}, \overline{\nu}) \xrightarrow{d} (\overline{M}, \overline{\nu} + d)$, then all places p_t are empty in $\overline{M}$, so that the move $(M, \nu) \xrightarrow{d} (M, \nu + d)$ is also possible in $\mathcal{N}$.

The cases for a transition $t \in T \setminus T'$ are straightforward. Thus $\mathcal{R}$ is a bisimulation. Assuming that $\mathcal{N}$ is bounded, the boundedness of $\overline{\mathcal{N}}$ is mainly due to the following observation: if a place p_t is unbounded in $\overline{\mathcal{N}}$ then any input place of t is unbounded in $\mathcal{N}$. $\square$

Proposition 2. *Let $\mathcal{N}$ be a time Petri net with atomic semantics. There exists a TPN $\overline{\mathcal{N}}$ with persistent atomic semantics which is weakly timed bisimilar to $\mathcal{N}$. The size of $\overline{\mathcal{N}}$ is linear w.r.t. the size of $\mathcal{N}$. Furthermore if $\mathcal{N}$ is bounded then $\overline{\mathcal{N}}$ is bounded.*

Proof. Here again, the construction is simple. Note that the only difference between the two semantics concerns the question whether a transition t can newly enable itself. With atomic semantics, this is the case as soon as t is enabled in the new marking while with persistent atomic semantics, this is never possible. In order to ensure that a transition t will be newly enabled if it is enabled in the new marking, we add an input place En_t^+ and an output place En_t^- to the

transition, with an instantaneous loop b_t leading back to En_t^+ , once the transition has been fired. The construction is represented in Figure 5, again from left to right.

Fig. 5. From A to PA

We consider the equivalence relation $\mathcal{R}$ which contains all pairs $((M, \nu), (\overline{M}, \overline{\nu}))$ such that:

- $M(p) = \overline{M}(p)$ for all places p in P , and
- for a transition $t \in En(M)$, $\nu(t) = \overline{\nu}(t)$ if t is enabled in $\overline{M}$ and 0 otherwise.

Again the latter case corresponds in $\mathcal{N}$ to a newly enabled transition.

Like in the previous proposition, the proof is mostly based on the fact that from any configuration $(\overline{M}, \overline{\nu})$, we can reach instantaneously a configuration $(\overline{M}_1, \overline{\nu}_1)$ such that $\overline{M}_1(En_t^+) = 1$ for all t , with the firing of a (possibly empty) sequence of transitions b_t , and again $(\overline{M}_1, \overline{\nu}_1)$ still equivalent to (M, ν) . Furthermore the claim about boundedness is mainly due to the following invariants $\overline{M}(En_t^+) + \overline{M}(En_t^-) = 1$. $\square$

3.2 Strict inclusion

The next proposition shows that the expressive power of TPNs depends on the chosen semantics even in the bounded case.

Proposition 3. *There exists a bounded TPN $\mathcal{N}$ with persistent atomic semantics such that no TPN (even unbounded) with atomic semantics is bisimilar to $\mathcal{N}$.*

Proof. Consider the following bounded TPN with PA semantics composed by a single transition t labeled by ε : $\text{ } \text{ } \text{ } t, \varepsilon, [0, 1]$

The (observable) behaviour of this net is simply to let the time elapse without reaching 1. Suppose that there is a TPN $\mathcal{N}$ with atomic semantics bisimilar to this TPN and let d_{min} be the minimum of the non null upper bounds occurring in the intervals associated with the transitions of $\mathcal{N}$ and 0.5 (in fact any value less than 1 would be convenient).

There must be a sequence $(M_0, \nu_0) \xrightarrow{d_0 t_1 \dots t_k d_k} (M, \nu)$ with $\sum_{i=0}^k d_i = 1 - d_{min}/2$ and (M, ν) bisimilar to $(\emptyset, 1 - d_{min}/2)$.

From (M, ν) , we fire or disable the transitions enabled at this configuration, which leads to a new configuration (M', ν') bisimilar to some $(\emptyset, 1 - \delta')$ with $0 < \delta' \leq d_{min}/2$. Now since (M', ν') is bisimilar to $(\emptyset, 1 - \delta')$ there must be a sequence $(M', \nu') \xrightarrow{d'_0 t'_1 \dots t'_{k'} d'_{k'}} \dots$ with $0 < \sum_{i=0}^{k'} d'_i < \delta'$.

Choose the first $d'_i > 0$ and let (M^*, ν^*) be the state reached before the duration d'_i . Since time may elapse in this state, all enabled transitions have non null upper bound for their interval, hence these bounds are greater than or equal to d_{min} . Since the transitions have been enabled at or after configuration (M', ν') , we have $\forall t, \nu^*(t) \leq d_{min}/2 - \delta' < d_{min}/2$, thus $(M^*, \nu^*) \xrightarrow{d_{min}/2} \dots$. But (M^*, ν^*) is bisimilar to $(\emptyset, 1 - \delta')$ which cannot let time elapse for a duration of $d_{min}/2$. This is a contradiction.

A similar proof could be developed with any interval $]a, b[$ or $[a, b[$ instead of $[0, 1[$.

□

4 Equivalence result for bounded TPNs with upper-closed intervals

Due to the strict inclusion established by proposition 3, we now restrict our study to bounded TPNs, with *upper-closed* intervals, *i.e.* with intervals $[a, b]$, $[a, \infty[$, $]a, b]$ or $]a, \infty[$.

Proposition 4. *Let $\mathcal{N}$ be a bounded TPN with upper-closed intervals and with persistent atomic semantics. There exists a bounded TPN $\overline{\mathcal{N}}$ with intermediate semantics which is weakly timed bisimilar to $\mathcal{N}$. The size of $\overline{\mathcal{N}}$ is linear w.r.t. the size of $\mathcal{N}$ and the logarithm of the bound.*

Proof. In this case, the construction of $\overline{\mathcal{N}}$ is more involved. Like above, we show how to simulate a transition t equipped with interval $[a, b]$, for $a \leq b$ (the other cases are similar). We first build a time subnet for t (Figure 6 below), to simulate time elapsing from the last time t was enabled, until reaching (and staying inside) interval $[a, b]$. The token is in place $start_t$ if the transition is enabled in the initial marking. The double arrow at the end indicates that the place end_t is both an input and an output place for the corresponding transition: time cannot progress. Clearly this part of the construction does not apply when the interval is not upper-closed.

Now, using the fact that the TPN is bounded, we consider its upper bound B and we associate with each place p a complementary place $\overline{p}$ such that for any reachable marking M , $M(\overline{p}) = B - M(p)$. Figure 7 represents a part of the subnet (on the right) for transition t (on the left), where $test_1$ is the beginning of the test step for what timing updates are required by the firing of t , and *Mutex* ensures that the updates are done (instantaneously) before anything else, as explained further.

Fig. 6. Time subnet for interval $I(t) = [a, b]$ of transition t

Fig. 7. From a transition in $\mathcal{N}$ to its translation in $\overline{\mathcal{N}}$

The remaining part of $\overline{\mathcal{N}}$ is devoted to the test of the other transitions from the original TPN, including t itself. Consider a given transition (say t_i), with again two input places p_1^i and p_2^i . The corresponding subnet consists of 4 modules, one for each case, depending on whether t_i can be fired or not before and after t . For this, two additional places are associated with t_i : E_{t_i} , which contains a token if t_i was enabled before the firing of t and NE_{t_i} its complementary place. If t_i is initially enabled then E_{t_i} is initially marked otherwise NE_{t_i} is marked. This group of 4 modules has a common input place $test_i$ and a common output place $test_{i+1}$, which means that the tests are to be executed sequentially (and instantaneously), except for the last one where all outgoing transitions are linked to $Mutex$. These places are not shown in the following figures.

Case 1: transition t_i is enabled both before and after firing t . To test this case, we use the simple module on the left of Figure 8, where E_{t_i} (test before t) and p_1^i and p_2^i (test after t) are input and output places.

Case 2: t_i is not enabled before but enabled after firing t . The module is very similar to the previous one and is on the right of Figure 8. Note that, in this case only, because of the PA semantics, there must be a reset on the valuation of the transition, which explains why the initial place $start_{t_i}$ of the time module for t_i is an output place.

Case 3: t_i is enabled neither before nor after firing t . To test this, we must find an input place of t_i , where the current number of tokens disable t_i . Here is the point where the boundedness hypothesis is required. In order to perform this test, we check whether $B - \bullet t_i(p) + 1$ tokens can be removed from a complementary place $\overline{p}$.

Fig. 8. Testing transition t_i : cases 1 and 2

Fig. 9. Testing transition t_i : case 3

Case 4: t_i is enabled before but not after firing t . In this case, we have a module (see Figure 10) similar to the one above, except that we must also test for all the different configurations of the time subnet corresponding to t_i , to disable the transitions by removing the tokens.

It can be seen in Figure 10 that there is a transition for each pair $(\bar{p}, state)$, where p is an input place of t_i and $state$ may be either the place $start_{t_i}$, the pair of places $(in_{t_i}, true_{[a_i, b_i]})$ or the pair $(end_{t_i}, true_{[a_i, b_i]})$. Like above, an edge from $\bar{p}$ to a transition must be labeled with $B - \bullet t_i(p) + 1$ (which is omitted in the figure). Note that the number of states (here 3) depends on the kind of intervals but is bounded by a constant. Thus this construction is still linear w.r.t. the number of input arcs of t .

We consider the equivalence relation $\mathcal{R}$ containing all pairs $((M, \nu), (\bar{M}, \bar{\nu}))$ such that

- M in $\mathcal{N}$ is obtained by projection: $M(p) = \bar{M}(p)$ for each place $p \in P$,
- for a transition t in T enabled by M : $\nu(t) = 0$ if the time subnet of t is empty, $\nu(t) = \bar{\nu}(l_t)$ if the place $start_t$ contains a token, $\nu(t) = a + \bar{\nu}(u_t)$ if the place in_t contains a token and $\nu(t) = b$ if the place end_t contains a token. Note that in both latter cases, $true_{[a, b]}$ also contains a token and the transition t can be fired in $\bar{\mathcal{N}}$.

Also note that if $\bar{M}(start_t) = 1$ and $\bar{\nu}(u_t) = a$, then with instantaneously firing u_t , transition t can also be fired. By a development similar to the previous ones, we can show that $\mathcal{R}$ is a bisimulation relation. More precisely, the proof is mainly based on emptying sequences from a configuration $(\bar{M}, \bar{\nu})$ of $\bar{\mathcal{N}}$: it is always possible to reach instantaneously a configuration $(\bar{M}_1, \bar{\nu}_1)$ such that the

Fig. 10. Testing transition t_i : case 4 (names of transitions are omitted)

testing subnet is empty, with $(\bar{M}_1, \bar{\nu}_1)$ still equivalent to (M, ν) . The details are omitted. $\square$

We conclude this section with the following table summarizing our results.

(Bounded) TPNs general intervals	I semantics	A semantics	PA semantics
I semantics		$\leq_{\mathcal{W}}$	$<_{\mathcal{W}}$
A semantics			$<_{\mathcal{W}}$
Bounded TPNs upper-closed intervals	I semantics	A semantics	PA semantics
I semantics		$\approx_{\mathcal{W}}$	$\approx_{\mathcal{W}}$

5 Conclusion

Since the introduction of TPNs, numerous works have lead to verification algorithms, tool development and applications to real-time systems. In this paper, we investigated some questions relative to the expressiveness of three different semantics for TPNs. First, we presented some standard patterns of discrete event systems where the original semantics (i.e. the intermediate one) is not the most appropriate, thus showing that atomic and persistent atomics semantics could be interesting alternatives.

Then we undertook a theoretical analysis of the three semantics w.r.t. weak time bisimilarity. We established a hierarchy between these semantics and proved

that the PA semantics is strictly more expressive than the two other ones. Since the last result is due to non upper-closed intervals and is still valid for bounded nets, we focused our study on bounded nets with upper-closed intervals. In this last case, the three semantics are equivalent.

However, even in the restricted case we feel that PA semantics should be preferred to the other ones since the emulation of the other semantics by the PA semantics is natural whereas the reverse emulation is rather involved and yields a TPN whose readability is doubtful. Furthermore most verification techniques are based on the class graph construction which applies with the same complexity for these three semantics.

References

1. P. A. Abdulla and A. Nylén. Timed Petri nets and BQOs. In *ICATPN'01*, volume 2075 of *LNCS*, pages 53–72. Springer-Verlag, june 2001.
2. M. Ajmone Marsan and G. Balbo and G. Conte and S. Donatelli and G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets*. Wiley Series in Parallel Computing, John Wiley and Sons, 1994, ISBN: 0-471-93059-8.
3. R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science B*, 126:183–235, 1994.
4. T. Aura and J. Lilius. A causal semantics for time Petri nets. *Theoretical Computer Science*, 243(1–2):409–447, 2000.
5. B. Berthomieu and M. Diaz. Modeling and verification of time dependent systems using time Petri nets. *IEEE Transactions on Software Engineering*, 17(3):259–273, March 1991.
6. Bernard Berthomieu and Pierre-Olivier Ribet and François Vernadat. The tool TINA - Construction of Abstract State Spaces for Petri Nets and Time Petri Nets International Journal of Production Research, Vol.42, N°14, pp.2741-2756, 2004.
7. B. Berthomieu and F. Vernadat. State class constructions for branching analysis of time Petri nets. In *TACAS'2003*, volume 2619 of *LNCS*, pages 442–457, 2003.
8. Franck Cassez and Olivier H. Roux. Structural Translation of Time Petri Nets into Timed Automata. In Michael Huth, editor, *Workshop on Automated Verification of Critical Systems (AVoCS'04)*, Electronic Notes in Computer Science. Elsevier, August 2004.
9. D. de Frutos Escrig, V. Valero Ruiz, and O. Marroquín Alonso. Decidability of properties of timed-arc Petri nets. In *ICATPN'00*, Aarhus, Denmark, volume 1825 of *LNCS*, pages 187–206, june 2000.
10. M. Diaz and P. Senac. Time stream Petri nets: a model for timed multimedia information. In *ATPN'94*, volume 815 of *LNCS*, pages 219–238, 1994.
11. G. Gardey, D. Lime, and O. (H.) Roux. ROMÉO: A tool for Time Petri Nets Analysis, 2003. The tool can be freely downloaded from www.irccyn.ec-nantes.fr/irccyn/d/fr/equipements/TempsReel/logs.
12. W. Khansa, J.P. Denat, and S. Collart-Dutilleul. P-Time Petri Nets for manufacturing systems. In *WODES'96*, Scotland, pages 94–102, 1996.
13. J. Lilius. Efficient state space search for time Petri nets. *Electronic Notes in Theoretical Computer Science*, volume 18, 1999.
14. D. Lime and O. H. Roux. State class timed automaton of a time Petri net. In *PNPM'03*. IEEE Computer Society, September 2003.

15. P. M. Merlin. *A study of the recoverability of computing systems*. PhD thesis, University of California, Irvine, CA, 1974.
16. L. Popova. On time Petri nets. *Journal Information Processing and Cybernetics, EIK*, 27(4):227–244, 1991.
17. C. Ramchandani. *Analysis of asynchronous concurrent systems by timed Petri nets*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1974.