



HAL
open science

Weak Time Petri Nets strike back!

Pierre-Alain Reynier, Arnaud Sangnier

► **To cite this version:**

| Pierre-Alain Reynier, Arnaud Sangnier. Weak Time Petri Nets strike back!. 2009. hal-00374482v3

HAL Id: hal-00374482

<https://hal.science/hal-00374482v3>

Preprint submitted on 17 Jun 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.

Weak Time Petri Nets strike back!

Pierre-Alain Reynier^{1,*} and Arnaud Sangnier^{2,**}

¹ LIF, Université Aix-Marseille & CNRS, France
pierre-alain.reynier@lif.univ-mrs.fr

² Dipartimento di Informatica, Università di Torino, Italy
sangnier@di.unito.it

Abstract. We consider the model of Time Petri Nets where time is associated with transitions. Two semantics for time elapsing can be considered: the strong one, for which all transitions are urgent, and the weak one, for which time can elapse arbitrarily. It is well known that many verification problems such as the marking reachability are undecidable with the strong semantics. In this paper, we focus on Time Petri Nets with weak semantics equipped with three different memory policies for the firing of transitions. We prove that the reachability problem is decidable for the most common memory policy (intermediate) and becomes undecidable otherwise. Moreover, we study the relative expressiveness of these memory policies and obtain partial results.

1 Introduction

For verification purpose, *e.g.* in the development of embedded platforms, there is an obvious need for considering time features and the study of timed models has thus become increasingly important. For distributed systems, different timed extensions of Petri nets have been proposed which allow the combination of an unbounded discrete structure with dense-time variables.

There are several ways to express urgency in timed systems, as discussed in [16]. In timed extensions of Petri nets, two types of semantics are considered for time elapsing. In the *weak* semantics, all time delays are allowed whereas in the *strong* one, all transitions are urgent, *i.e.* time delays cannot disable transitions. While for models with finite discrete structure (such as timed extensions of bounded Petri nets or timed automata [3]), standard verification problems are decidable for both semantics, for models with infinite discrete structure, the choice of the semantics has a deep influence on decidability issues. In this work, we consider the model of Time Petri Nets [14] (TPN) where clocks are associated with transitions, and which is commonly considered under a strong semantics. In this model, all the standard verification problems are known to be undecidable [10]. On the other hand, in the model of timed-arc Petri nets [5], where clocks are associated with tokens and which is equipped with a weak semantics, many verification problems are decidable (coverability, boundedness...). Indeed, this

* Partly supported by the French project DOTS (ANR-06-SETI-003) and by the European project QUASIMODO (FP7-ICT-STREP-214755).

** Supported by a post-doctoral scholarship from DGA/ENS Cachan and by the MIUR-PRIN project PaCo - Performability-Aware Computing: Logics, Models and Languages.

semantics entails for this model monotonicity properties which allow the application of well-quasi-ordering techniques, see [8, 2, 1]. Note however that the reachability of a discrete marking is undecidable, as proven in [17]. A natural question, which had surprisingly no answer until now, as mentioned in a recent survey on the topic [7], is thus to study TPN under a weak semantics of time elapsing.

The time-elapsing policy states which delays are allowed in a configuration. The memory policy is concerned with the resets of clocks, and intuitively specifies, when firing a transition, which timing informations are preserved. The original model of Merlin [14] is equipped with an *intermediate* semantics which considers the intermediary marking between consumption and production. Two other memory policies have been considered in [4] (the *atomic* and the *persistent atomic*) in which the firings of transitions are performed atomically. While these policies can be thought as cosmetic for the model of TPN, the results we obtain show this is not the case.

We are interested in the impact of the weak semantics on TPN, distinguishing between the different memory policies. We first study the decidability issues, and prove that for TPN with weak intermediate semantics, a discrete marking is reachable if and only if it is reachable in the underlying untimed Petri net. As a corollary, the problem of the marking reachability (and also coverability, boundedness) is decidable for this model. More surprisingly, we also prove that when changing the memory policy this result does not hold anymore and the verification problems become undecidable. In this work, we only consider untimed verification problems and we plan to study timed versions in future work. We then compare w.r.t. weak time bisimilarity (weak stands here for silent transitions) the expressive power of weak TPN looking at the different memory policies. We first prove that the persistent atomic semantics is strictly more expressive than the atomic semantics. Then, concerning atomic and intermediate memory policies, we provide a TPN which shows that the atomic semantics is not included in the intermediate one.

Related works. As mentioned above, there are, up to our knowledge, only very few works considering TPN under a weak semantics. In [7] the authors have proven that the weak intermediate semantics and the strong intermediate semantics are incomparable. In another line of work, [9] considers TPN under a semantics which is a kind of compromise between the standard strong and weak semantics. They provide translations between this model and timed state machines.

2 Definitions

Let Σ be a finite alphabet, Σ^* is the set of finite words over Σ . We note $\Sigma_\tau = \Sigma \cup \{\tau\}$ where $\tau \notin \Sigma$ represents internal actions. ϵ will represent the empty word. The sets \mathbb{N} , \mathbb{Q} , $\mathbb{Q}_{\geq 0}$, \mathbb{R} and $\mathbb{R}_{\geq 0}$ are respectively the sets of natural, rational, non-negative rational, real and non-negative real numbers. A *valuation* v over a finite set X is a mapping in $\mathbb{R}_{\geq 0}^X$. For $v \in \mathbb{R}_{\geq 0}^X$ and $d \in \mathbb{R}_{\geq 0}$, $v+d$ denotes the valuation defined by $(v+d)(x) = v(x)+d$. We note $\mathbf{0}$ the valuation which assigns to every $x \in X$ the value 0.

As commonly in use for Time Petri Nets, we will associate rational intervals with transitions. Note that we could handle intervals with bounds given as real numbers if we abstract the problem of comparison of real numbers. We consider the set $\mathcal{I}(\mathbb{Q}_{\geq 0})$ of

non-empty intervals (a, b) with non-negative rational bounds $a, b \in \mathbb{Q}_{\geq 0}$. We consider both open and closed bounds, and also allow a right open infinite bound as in $[2, +\infty[$.

2.1 Petri Nets

Definition 1 (Labeled Petri Net (PN)). A Labeled Petri Net over the alphabet Σ_τ is a tuple $(P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda)$ 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,
- $\Lambda : T \rightarrow \Sigma_\tau$ is the labeling function

As commonly in use in the literature, the vector $\bullet(t)$ (resp. $(t)^\bullet$) in \mathbb{N}^P is noted $\bullet t$ (resp. t^\bullet). The semantics of a PN $\mathcal{N} = (P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda)$ is given by its associated labeled transition system $\llbracket \mathcal{N} \rrbracket = (\mathbb{N}^P, M_0, \Sigma_\tau, \Rightarrow)$ where $\Rightarrow \subseteq \mathbb{N}^P \times \Sigma_\tau \times \mathbb{N}^P$ is the transition relation defined by $M \xrightarrow{a} M'$ iff $\exists t \in T$ s.t. $\Lambda(t) = a \wedge M \geq \bullet t \wedge M' = M - \bullet t + t^\bullet$. For convenience we will sometimes also write, for $t \in T$, $M \xrightarrow{t} M'$ if $M \geq \bullet t$ and $M' = M - \bullet t + t^\bullet$. We also write $M \Rightarrow M'$ if there exists $a \in \Sigma_\tau$ such that $M \xrightarrow{a} M'$. The relation \Rightarrow^* represents the reflexive and transitive closure of \Rightarrow . We denote by $\text{Reach}(\mathcal{N})$ the set of reachable markings defined by $\{M \in \mathbb{N}^P \mid M_0 \Rightarrow^* M\}$.

It is well known that for PN the reachability problem which consists in determining whether a given marking M belongs to $\text{Reach}(\mathcal{N})$ is decidable; it has in fact been proved independently in [13] and [12].

We introduce a last notation concerning Labeled Petri Nets. Given a PN \mathcal{N} , a marking M of \mathcal{N} and a multi-set $\Delta = \langle t_1, \dots, t_n \rangle$ of transitions of \mathcal{N} , we write $M \xrightarrow{\Delta} M'$ if and only if the multi-set Δ can be fired from M , meaning that there exists an ordering of transitions in Δ , represented as a permutation φ of $\{1, \dots, n\}$, such that the sequence of firings $M \xrightarrow{t_{\varphi(1)}} M_1 \xrightarrow{t_{\varphi(2)}} M_2 \dots \xrightarrow{t_{\varphi(n)}} M'$ exists in $\llbracket \mathcal{N} \rrbracket$.

2.2 Timed Transition Systems

Timed transition systems describe systems which combine discrete and continuous evolutions. They are used to define the behavior of timed systems such as Time Petri Nets [14] or Timed Automata [3].

Definition 2 (Timed Transition System (TTS)). A timed transition system over the alphabet Σ_τ is a transition system $S = (Q, q_0, \Sigma_\tau, \rightarrow)$, where the transition relation $\rightarrow \subseteq Q \times (\Sigma_\tau \cup \mathbb{R}_{\geq 0}) \times Q$ consists of discrete transitions $q \xrightarrow{a} q'$ (with $a \in \Sigma_\tau$) representing an instantaneous action, and continuous transitions $q \xrightarrow{d} q'$ (with $d \in \mathbb{R}_{\geq 0}$) representing the passage of d units of time.

Moreover, we require the following standard properties for 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 $q \xrightarrow{d'} q'' \xrightarrow{d''} q'$.

With these properties, a *run* of S can be defined as a finite sequence of moves $\rho = q_0 \xrightarrow{d_0} q'_0 \xrightarrow{a_0} q_1 \xrightarrow{d_1} q'_1 \xrightarrow{a_1} q_2 \dots \xrightarrow{a_n} q_{n+1}$ where discrete and continuous transitions alternate. To such a run corresponds the timed word $w = (a_i, \eta_i)_{0 \leq i \leq n}$ over Σ_τ where $\eta_i = \sum_{j=0}^i d_j$ is the time at which a_i happens. We then denote by $Untimed(w)$ the projection of the word $a_0 a_1 \dots a_n$ over the alphabet Σ and by $Duration(w)$ the duration η_n . Note that in the word $Untimed(w)$ the symbol τ does not appear. We will sometimes apply, without possible ambiguities, these notations to runs writing $Untimed(\rho)$ and $Duration(\rho)$. We might also describe the run writing directly $q_0 \xrightarrow{w} q_{n+1}$.

2.3 Time Petri Nets

Syntax. Introduced in [14], Time Petri Nets associate a time interval with each transition of a Petri net.

Definition 3 (Labeled Time Petri Net (TPN)). A Labeled Time Petri Net over the alphabet Σ_τ is a tuple $(P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda, I)$ where:

- $(P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda)$ is a PN,
- $I : T \mapsto \mathcal{I}(\mathbb{Q}_{\geq 0})$ associates with each transition a firing interval.

In the sequel, we associate with an interval its left bound and its right bound. More generally, given a transition t of a TPN, we will denote by $eft(t)$ (resp. $lft(t)$) the left bound of $I(t)$ (resp. the right bound of $I(t)$), standing for earliest firing time (resp. latest firing time). We have hence $I(t) = (eft(t), lft(t))$.

Semantics. A *configuration* of a TPN is a pair (M, ν) , where M is a *marking* over P , i.e. a mapping in \mathbb{N}^P , with $M(p)$ the number of tokens in place p . A transition t is *enabled* in a marking M if $M \geq \bullet t$. We denote by $En(M)$ the set of enabled transitions in M . The second component of the pair (M, ν) is a *valuation* over $En(M)$, i.e. a mapping in $\mathbb{R}_{\geq 0}^{En(M)}$. Intuitively, for each enabled transition t in M , $\nu(t)$ represents the amount of time that has elapsed since t is enabled. An enabled transition t can be fired if $\nu(t)$ belongs to the interval $I(t)$. The marking obtained after this firing is as usual the new marking $M' = M - \bullet t + t^\bullet$. Moreover, some valuations are reset and we say that the corresponding transitions are newly enabled.

Different semantics can be chosen in order to realize these resets. This choice depends of what is called the *memory policy*. For $M \in \mathbb{N}^P$ and $t, t' \in T$ such that $t \in En(M)$ we define in different matters a predicate $\uparrow enabled_s(t', M, t)$ with $s \in \{I, A, PA\}$ which is true if t' is *newly enabled* by the firing of transition t from marking M , and false otherwise. This predicate indicates whether we need to reset the clock of t' after firing the transition t at the marking M .

I: The *intermediate semantics* considers that the firing of a transition is performed in two steps: consuming the tokens in $\bullet t$, and then producing the tokens in t^\bullet . Intuitively, it resets the clocks of t and of the transitions that could not be fired in parallel with t from the marking M . Formally, the predicate $\uparrow enabled_I(t', M, t)$ is defined by:

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

A: The *atomic semantics* considers that the firing of a transition is obtained by an atomic step. It resets the clocks of t and of the transitions t' which are not enabled at M . The corresponding predicate $\uparrow enabled_A(t', M, t)$ is defined by:

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

PA: The *persistent atomic semantics* behaves as the atomic semantics except that it does not reset the clock of t .

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

Finally, as recalled in the introduction, there are two ways of letting the time elapse in TPN. The first way, known as the *strong semantics*, is defined in such a matter that time elapsing cannot disable a transition. Hence, when the upper bound of a firing interval is reached then the transition must be fired or disabled. In contrast to that the *weak semantics* does not make any restriction on the elapsing of time. In this work, we focus on the weak semantics of TPN.

Definition 4 (Weak semantics of a TPN). Let $s \in \{I, A, PA\}$. The weak s -semantics of a TPN $\mathcal{N} = (P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda, I)$ is a timed transition system $\llbracket \mathcal{N} \rrbracket_s = (Q, q_0, \Sigma_\tau, \rightarrow_s)$ where $Q = \mathbb{N}^P \times \mathbb{R}_{\geq 0}^{En(M)}$, $q_0 = (M_0, \mathbf{0})$ and \rightarrow_s consists of discrete and continuous moves:

– the discrete transition relation is defined $\forall a \in \Sigma_\tau$ by:

$$(M, \nu) \xrightarrow{a}_s (M', \nu') \text{ iff } \exists t \in T \text{ s.t. } \begin{cases} \Lambda(t) = a, \text{ and,} \\ t \in En(M) \wedge M' = M - \bullet t + t^\bullet, \text{ and,} \\ \nu(t) \in I(t), \text{ and,} \\ \forall t' \in En(M'), \\ \nu'(t') = \begin{cases} 0 & \text{if } \uparrow enabled_s(t', M, t) \\ \nu(t') & \text{otherwise} \end{cases} \end{cases}$$

– the continuous transition relation is defined $\forall d \in \mathbb{R}_{\geq 0}$ by:

$$(M, \nu) \xrightarrow{d}_s (M, \nu') \text{ iff } \nu' = \nu + d$$

We also write a discrete transition $(M, \nu) \xrightarrow{t}_s (M', \nu')$ to characterize the transition $t \in T$ which allows the firing $(M, \nu) \xrightarrow{\Lambda(t)}_s (M', \nu')$. We extend this notation to words $\theta \in (T \cup \mathbb{R}_{\geq 0})^*$, which correspond to sequences of transitions and delays and

lead to a unique (if it exists) run ρ . We may write this run $\rho : (M, \nu) \xrightarrow{\theta}_s (M', \nu')$ and use $Untimed(\theta)$ (resp. $Duration(\theta)$) to denote the word $Untimed(\rho)$ (resp. to represent the delay $Duration(\rho)$). Finally, for $s \in \{I, A, PA\}$, we write $(M, \nu) \rightarrow_s (M', \nu')$ if there exists $a \in \Sigma_\tau \cup \mathbb{R}_{\geq 0}$ such that $(M, \nu) \xrightarrow{a}_s (M', \nu')$. The relation \rightarrow_s^* denotes the reflexive and transitive closure of \rightarrow_s . For a TPN \mathcal{N} with an initial marking M_0 we define the following reachability sets according to the considered semantics: $\text{Reach}(\mathcal{N})_s = \{(M, \nu) \mid (M_0, \mathbf{0}) \rightarrow_s^* (M, \nu)\}$.

Example 1. We illustrate the impact of the three memory policies in weak semantics.




Fig. 1. The TPN \mathcal{N}_1 .

Consider the net depicted on Figure 1, and the execution $(M, \mathbf{0}) \xrightarrow{1}_s (M, \mathbf{1}) \xrightarrow{a}_s (M, \nu)$ where $M(p) = 1$. With the intermediate semantics, both clocks are reset as in the intermediate marking, the place p is empty. With the atomic semantics, the clock associated with $t2$ is not reset and the clock associated with $t1$ is reset because it corresponds to the fired transition. Finally, with the persistent atomic semantics no clock is reset.

3 Decidability

3.1 Considered problems and known results

Assume $\mathcal{N} = (P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda, I)$ is a TPN. In this section, we will consider the following problems for $s \in \{I, A, PA\}$:

- (1) The *marking reachability problem* : given $M \in \mathbb{N}^P$, does there exist $\nu \in \mathbb{R}_{\geq 0}^{En(M)}$ such that $(M, \nu) \in \text{Reach}(\mathcal{N})_s$?
- (2) The *marking coverability problem* : given $M \in \mathbb{N}^P$, does there exist $M' \in \mathbb{N}^P$ and $\nu \in \mathbb{R}_{\geq 0}^{En(M')}$ such that $M' \geq M$ and $(M', \nu) \in \text{Reach}(\mathcal{N})_s$?
- (3) The *boundedness problem* : does there exist $b \in \mathbb{N}$ such that for all $(M, \nu) \in \text{Reach}(\mathcal{N})_s$ and for all $p \in P$, $M(p) \leq b$?

It is well known that the "untimed" versions of these problems are decidable in the case of Petri nets. In fact, as mentioned before the marking reachability problem is decidable for Petri nets [12, 13] and the two other problems can be solved using the Karp and Miller tree whose construction is given in [11].

From [10], we know that these problems are all undecidable when considering TPN with strong semantics no matter whether the semantics is intermediate, atomic or persistent atomic. In fact a TPN with strong semantics can simulate a Minsky machine. A Minsky machine manipulates two integer variables c_1 and c_2 and is composed of a finite number of instructions, each of these instructions being either an incrementation ($q : c_i := c_i + 1$) or a decrementation with a test to zero ($q : \text{if } c_i = 0 \text{ goto } q' \text{ else } c_i := c_i - 1; \text{ goto } q''$), where $i \in \{1, 2\}$ and q, q', q'' are some labels preceding each instruction. There is also a special label q_f from which the machine cannot do anything.

In [15], Minsky proved that the halting problem, which consists in determining if the instruction labeled with q_f is reachable, is undecidable.

It is easy to encode an incrementation using a TPN (or even a PN), with a transition consuming a token in a place characterizing the current control state and producing a token in the next control state and in a place representing the incremented counter.

When encoding the decrementation with the test to zero, the strong semantics plays a crucial role. This encoding is represented on Figure 2. If there is a token in the place c_i , there is no way for the TPN to produce a token in the place q' because time cannot elapse since the transition labeled with the interval $[0, 0]$ is firable. The example of the Figure 2 shows that the strong time semantics allows to encode priorities (between transitions in conflict) and thus to encode inhibitor arcs. This construction obviously fails with the weak semantics.




Fig. 2. Encoding decrementation with strong semantics.

3.2 The peculiar case of TPN with weak intermediate semantics

We prove here that the undecidability results we had before in the case of TPN with strong semantics do not hold anymore when considering the weak intermediate semantics. Before proving this we introduce some notations. For a TPN $\mathcal{N} = (P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \lambda, I)$, we denote by \mathcal{N}^U the untimed PN obtained by removing from \mathcal{N} the component I . Furthermore given a set of configurations $C \subseteq \mathbb{N}^P \times \mathbb{R}_{\geq 0}^T$ of \mathcal{N} , we denote by $\text{Untime}(C)$ the projection of C over the set \mathbb{N}^P . For $s \in \{I, A, PA\}$, we have by definition of the different semantics that $\text{Untime}(\text{Reach}(\mathcal{N})_s) \subseteq \text{Reach}(\mathcal{N}^U)$ and as shown by the example given in Figure 2 this inclusion is strict in the case of the strong semantics. When considering the **weak intermediate** semantics, we prove that from any sequence of transitions Δ firable in $\llbracket \mathcal{N}^U \rrbracket$, we can effectively compute a reordering of Δ , and associated timestamps, leading to a correct run of $\llbracket \mathcal{N} \rrbracket_I$.

Theorem 5. For all TPN \mathcal{N} , $\text{Untime}(\text{Reach}(\mathcal{N})_I) = \text{Reach}(\mathcal{N}^U)$.

According to the previous remark, we only have to prove that $\text{Reach}(\mathcal{N}^U) \subseteq \text{Untime}(\text{Reach}(\mathcal{N})_I)$. Therefore, we first state the following property expressing that if we reduce the intervals associated with transitions, this restricts the set of reachable configurations:

Lemma 6. Let $\mathcal{N}, \mathcal{N}'$ be two TPN identical except on their last component associating intervals to transitions, say respectively I and I' . If we have $I'(t) \subseteq I(t)$ for any $t \in T$, then $\text{Reach}(\mathcal{N}')_I \subseteq \text{Reach}(\mathcal{N})_I$.

In the sequel, we will consider TPN in which intervals are reduced to singletons. That is we have $I(t) = [eft(t), lft(t)]$ with $eft(t) = lft(t)$ for all transitions $t \in T$. The proof of the result in this particular case thus entails the result in the general case. Before to proceed we introduce additional definitions for TPN.

Given a TPN \mathcal{N} , a marking M of \mathcal{N} and Δ a multiset of transitions of \mathcal{N} , we define the set $Candidate(M, \Delta) = \{t \in \Delta \mid M \xrightarrow{t} M' \xRightarrow{\Delta \setminus t}\}$. We will then say that a configuration (M, ν) is *compatible* with a multiset Δ iff:

$$M \xRightarrow{\Delta} \text{ and } \forall t \in Candidate(M, \Delta), \nu(t) \leq lft(t).$$

We now prove the following proposition, which intuitively states how to turn a sequence of transitions in the untimed Petri net into a timed execution of the TPN.

Proposition 7. *Let \mathcal{N} be a TPN with singleton intervals and (M, ν) be a configuration of \mathcal{N} compatible with some multiset of transitions Δ . Then, for any transition $t \in Candidate(M, \Delta)$ such that $\delta(t) = lft(t) - \nu(t)$ is minimal (among the transitions of $Candidate(M, \Delta)$), we have:*

- (i) $(M, \nu) \xrightarrow{\delta(t)}_I (M, \nu + \delta(t)) \xrightarrow{t}_I (M', \nu')$,
- (ii) (M', ν') is compatible with $\Delta' = \Delta \setminus t$,

Proof. Let $t \in Candidate(M, \Delta)$ be such that for all $t' \in Candidate(M, \Delta)$, we have $lft(t) - \nu(t) = \delta(t) \leq \delta(t') = lft(t') - \nu(t')$.

- (i) First the time elapsing transition $(M, \nu) \xrightarrow{\delta(t)}_I (M, \nu + \delta(t))$ is possible as we consider the weak semantics. Second, the discrete transition $(M, \nu + \delta(t)) \xrightarrow{t}_I (M', \nu')$ is also possible as $\nu(t) + \delta(t) = lft(t)$ by definition of $\delta(t)$, and since the intervals associated with transitions are all singletons.
- (ii) To prove compatibility, first note that $M' \xRightarrow{\Delta'}$ because $t \in Candidate(M, \Delta)$. Second, let $t' \in Candidate(M', \Delta')$. We distinguish two cases according to the value of the predicate $\uparrow enabled_I(t, M, t')$:
 - If $\uparrow enabled_I(t, M, t')$ is true, then we have $\nu'(t') = 0$ and the result follows.
 - Otherwise, the definition of $\uparrow enabled_I(t, M, t')$ implies that $M \geq \bullet t + \bullet t'$. As a consequence, we have $M \xrightarrow{t'} \xrightarrow{t}$. Then as $t' \in Candidate(M', \Delta \setminus t)$ we get that $t' \in Candidate(M, \Delta)$. Due to the minimality of $\delta(t)$ among the set $Candidate(M, \Delta)$, we obtain $\nu'(t') = \nu(t') + \delta(t) \leq \nu(t') + \delta(t') = lft(t')$ as desired.

This concludes the proof. □

The inclusion $Reach(\mathcal{N}^U) \subseteq \text{Untime}(Reach(\mathcal{N})_I)$ in the case of TPN with singleton intervals easily follows from this result. Indeed, consider some reachable marking M in $Reach(\mathcal{N}^U)$. There exists a sequence of transitions that leads to M from M_0 , we represent it through some multiset Δ . As initially all clock valuations are null in $\llbracket \mathcal{N} \rrbracket_I$, the configuration $(M_0, \mathbf{0})$ is thus compatible with Δ . An induction on the size of Δ , together with Proposition 7, thus gives the result. Note that Proposition 7 describes an effective procedure to compute a timed execution of $\llbracket \mathcal{N} \rrbracket_I$: simply consider the transitions that are candidates, and choose one with the earliest deadline.

Using the decidability results in the case of PN, we obtain the following corollary:

Corollary 8. *The marking reachability, marking coverability and boundedness problems are decidable in the case of TPN with weak intermediate semantics.*

3.3 Undecidability for weak atomic and weak persistent atomic semantics

We consider now the case of the weak atomic and weak persistent atomic semantics. As for the strong semantics, but with a more involved construction, we will show that it is possible to encode the behavior of a Minsky machine into a TPN with weak (persistent) atomic semantics from which we will deduce the undecidability results. The TPN we build contains a place for each counter c_i with $i \in \{1, 2\}$ and a place for each label q of the considered Minsky machine. Furthermore, when executing the net, we will preserve the invariant that there is a single place corresponding to a label q which is marked.




Fig. 3. Encoding decrementation with weak atomic or persistent atomic semantics.

Encoding an incrementation can be done as in the strong semantics. Figure 3 shows how to encode the instruction ($q : \text{if } c_i = 0 \text{ goto } q' \text{ else } c_i := c_i - 1; \text{ goto } q''$) using a TPN with weak atomic or persistent atomic semantics. We now explain the idea of this encoding. We consider the two following cases for the net shown in Figure 3:

1. Assume that the only place which contains a token is the place q , which means we are in the case where the value of c_i is equal to 0 (no token in place c_i). The net then can only fire the sequence of transitions $t1, t2, t3$ and then $t4$ and finally it reaches a configuration where the only marked place is q' .
2. Assume now that there is a token in place q and that there is at least one token in place c_i . We are in the case where the value of c_i is different of 0. We have the following sequence of transitions:
 - only the transition $t1$ is fireable, so the net fires it;
 - afterwards the transition $t2$ and the transition $t3$ are fireable. In fact, since we are considering weak semantics the deadline of $t3$ can be ignored thus making time passage in order to fire $t2$. Note that if the net chooses to fire $t3$, it will reach a deadlock state where no more transitions can be fired without having put a token in the place q' or q'' , therefore we assume that the transition $t2$ is first fired;
 - after having waiting one time unit and firing $t2$, the only transition which can be fired is $t5$. In fact since we are considering atomic (or persistent atomic) semantics, firing $t2$ does not make $t3$ newly enabled, whereas the weak intermediate semantics would have reset the clock associated to $t3$. So the net fires

- $t5$ consuming the token in $p2, p3$ and two tokens in c_i (at least one was present from the initial configuration and the first firing of $t2$ added another one);
- finally the net ends in a configuration with one token in q'' and the place c_i contains one token less than in the initial configuration.

The above construction allows to reduce the halting problem for Minsky machine to the marking coverability problem for weak (persistent) atomic semantics. From this we can also deduce the undecidability for the marking reachability and boundedness problems. Hence:

Theorem 9. *The marking reachability, marking coverability and boundedness problems are undecidable for TPN with weak atomic or weak persistent atomic semantics.*

In comparison with what occurs in the case of the strong semantics, this result is surprising, and it reveals the important role played by the memory policy when considering the weak semantics. Recall that as we have seen earlier, with the strong semantics, these problems are undecidable no matter which memory policy is chosen.

Finally, in the above construction, we can replace the edges between $p2$ and $t2$ by a read arc. Consequently, the considered problems are also undecidable for weak intermediate TPN with read arcs, unlike what happens for timed-arc Petri nets [6].

4 Expressiveness

4.1 Preliminaries

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

- for $d \in \mathbb{R}_{\geq 0}$, $q \xrightarrow{d} q'$ iff there is a run ρ in S such that $\rho = q \xrightarrow{w} q'$ and $Untimed(w) = \varepsilon$ and $Duration(w) = d$,
- for $a \in \Sigma$, $q \xrightarrow{a} q'$ iff there is a run ρ in S such that $\rho = q \xrightarrow{w} q'$ and $Untimed(w) = a$ and $Duration(w) = 0$.

This allows us to define the following notion:

Definition 10 (Weak Timed Bisimulation). *Let $S_1 = (Q_1, q_0^1, \Sigma_\tau, \rightarrow_1)$ and $S_2 = (Q_2, q_0^2, \Sigma_\tau, \rightarrow_2)$ be two TTS and \sim be a binary relation over $Q_1 \times Q_2$. \sim is a weak timed bisimulation between S_1 and S_2 if and only if:*

- $q_0^1 \sim q_0^2$, and,
- for $a \in \Sigma \cup \mathbb{R}_{\geq 0}$, if $q_1 \xrightarrow{a}_1 q'_1$ and if $q_1 \sim q_2$ then there exists $q'_2 \in Q_2$ such that $q_2 \xrightarrow{a}_2 q'_2$ and $q'_1 \sim q'_2$; conversely if $q_2 \xrightarrow{a}_2 q'_2$ and if $q_1 \sim q_2$ then there exists $q'_1 \in Q_1$ such that $q_1 \xrightarrow{a}_1 q'_1$ and $q'_1 \sim q'_2$.

Two TTS S_1 and S_2 are *weak timed bisimilar* if there exists a weak timed bisimulation between S_1 and S_2 . We then write $S_1 \approx S_2$.

Definition 11 (Expressiveness w.r.t. Weak Timed Bisimilarity). *The class \mathcal{C} of TTS is less expressive than \mathcal{C}' w.r.t. weak timed bisimilarity if for all TTS $S \in \mathcal{C}$ there is a TTS $S' \in \mathcal{C}'$ such that $S \approx S'$. We write $\mathcal{C} \sqsubseteq \mathcal{C}'$. If moreover there is a $S' \in \mathcal{C}'$ such that there is no $S \in \mathcal{C}$ with $S \approx S'$, then \mathcal{C} is strictly less expressive than \mathcal{C}' . We then write $\mathcal{C} \sqsubset \mathcal{C}'$.*

For $s \in \{I, A, PA\}$, we will denote by \mathcal{TPN}_s the class of TTS induced by TPN with s -semantics.

4.2 Atomic versus Persistent Atomic semantics

In [4], the authors prove that for TPN with strong semantics, the persistent atomic semantics is more expressive than the atomic semantics. We prove here that this result still holds in the case of the weak semantics. Intuitively, as it is shown on Figure 4, from a TPN with atomic semantics, we build another TPN in which we duplicate each transition. During an execution of this last TPN, at most one of the transitions obtained after duplication is enabled, and when it is fired it cannot be enabled again at the next step whereas the other one can. This trick allows us to simulate the atomic semantics with the persistent atomic one.




Fig. 4. From atomic to persistent atomic semantics.

Proposition 12. *For all TPN \mathcal{N} , we can build a TPN \mathcal{N}' such that $\llbracket \mathcal{N} \rrbracket_A \approx \llbracket \mathcal{N}' \rrbracket_{PA}$.*

Proof. Let $\mathcal{N} = (P, T, \Sigma_\tau, \bullet(\cdot), (\cdot)^\bullet, M_0, \Lambda, I)$ be a TPN over Σ_τ . Figure 4 represents the construction of the TPN \mathcal{N}' . Formally, its set of places P' is equal to $P \cup \{p_t^1, p_t^2 \mid t \in T\}$ and its set of transitions T' contains two copies t^1 and t^2 of each transition $t \in T$. These copies are connected in the same way as the transition t is in \mathcal{N} , plus additional edges to the places p_t^1 and p_t^2 , as depicted on Figure 4. Finally the initial marking of \mathcal{N}' is M'_0 such that for all $p \in P$, $M'_0(p) = M_0(p)$ and for all $t \in T$, $M'_0(p_t^1) = 1$ and $M'_0(p_t^2) = 0$.

We now consider the relation $\sim \subseteq (\mathbb{N}^P \times \mathbb{R}_{\geq 0}^T) \times (\mathbb{N}^{P'} \times \mathbb{R}_{\geq 0}^{T'})$ between the configurations of $\llbracket \mathcal{N} \rrbracket_A$ and the ones of $\llbracket \mathcal{N}' \rrbracket_{PA}$ defined by $(M, \nu) \sim (M', \nu')$ iff:

- for all $p \in P$, $M(p) = M'(p)$ and for all $t \in T$, $M'(p_t^1) + M'(p_t^2) = 1$,
- for all $t \in T$, for all $i \in \{1, 2\}$ if $t \in \text{En}(M)$ and $t^i \in \text{En}(M')$ then $\nu(t) = \nu'(t^i)$.

It is then easy to verify that the relation \sim is a weak timed bisimulation. \square

We will now prove that the inclusion we obtain in the above proposition is strict. But before, we address a technical point which we will use to delay some sequences of transitions in weak TPN.

Lemma 13. *Let $s \in \{I, A, PA\}$ and consider a TPN \mathcal{N} such that b is the smallest positive upper bound of the intervals of \mathcal{N} . Let ρ be a run in $\llbracket \mathcal{N} \rrbracket_s$ of the form $\rho : (M, \nu) \xrightarrow{\delta > 0}_s (M, \nu + \delta) \xrightarrow{t_1}_s \dots \xrightarrow{t_n}_s$, such that there exists a value $\eta \geq 0$ verifying:*

- (i) $\forall i \in \{1, \dots, n\}, t_i \in \text{En}(M) \Rightarrow \nu(t_i) \leq \eta$,
- (ii) $\eta + \delta < \frac{b}{2}$

Then the sequence $\rho' : (M, \nu) \xrightarrow{\delta + \frac{b}{2}}_s (M, \nu') \xrightarrow{t_1}_s \dots \xrightarrow{t_n}_s$ is firable in $\llbracket \mathcal{N} \rrbracket_s$.

Proof. We introduce the following notations describing the run ρ :

$$(M, \nu) \xrightarrow{\delta > 0}_s (M, \nu + \delta) = (M_1, \nu_1) \xrightarrow{t_1}_s (M_2, \nu_2) \xrightarrow{t_2}_s (M_3, \nu_3) \dots (M_n, \nu_n) \xrightarrow{t_n}_s (M_{n+1}, \nu_{n+1}).$$

We split the set of indices $\{1, \dots, n\}$ into two disjoint subsets, $I_{=0} = \{i \in \{1, \dots, n\} \mid \nu_i(t_i) = 0\}$, and $I_{\neq 0} = \{i \in \{1, \dots, n\} \mid \nu_i(t_i) \neq 0\}$. These two sets contain respectively the index of the transitions fired with a null (respectively non null) valuation.

We first show the following property:

$$\forall i \in I_{\neq 0}, 0 < \nu_i(t_i) < \frac{b}{2} \text{ and } [\nu_i(t_i), b[\subseteq I(t_i) \quad (1)$$

Consider the first part of this property. Note that these valuations are strictly positive by definition of $I_{\neq 0}$, thus proving the left inequality. Consider now the right inequality, and let $i \in I_{\neq 0}$. Since the sequence $t_1 \dots t_i$ is instantaneous and follows a non null delay step, t_i is never newly enabled during $t_1 \dots t_{i-1}$ (otherwise $\nu_i(t_i) = 0$). As a consequence, we have $\nu_i(t_i) = \nu(t_i) + \delta$, and by properties (i) and (ii) of ρ , we obtain $\nu_i(t_i) \leq \eta + \delta < \frac{b}{2}$, as desired.

By definition of b , and since t_i is firable from (M_i, ν_i) , the inclusion $[\nu_i(t_i), b[\subseteq I(t_i)$ holds for any $i \in I_{\neq 0}$.

We consider now the run $\rho' : (M, \nu) \xrightarrow{\delta + \frac{b}{2}}_s (M, \nu + \delta + \frac{b}{2}) = (M_1, \nu'_1) \xrightarrow{t_1}_s (M_2, \nu'_2) \xrightarrow{t_2}_s (M_3, \nu'_3) \dots (M_n, \nu'_n) \xrightarrow{t_n}_s (M_{n+1}, \nu'_{n+1})$. Note that the increasing of the delay is possible because of the weak semantics: in the strong one, the modification of the delay step could be impossible. To prove that this sequence is firable in $\llbracket \mathcal{N} \rrbracket_s$, we proceed by contradiction. Assume there exists a position $i \in \{1, \dots, n\}$ such that t_i is not firable from (M_i, ν'_i) , and pick the smallest position verifying this property. We distinguish two cases:

1. **If $i \in I_{=0}$.** Then we have $\nu_i(t_i) = 0$. Since the instantaneous sequence $t_1 \dots t_i$ immediately follows in ρ the delay step $\delta > 0$, t_i is newly enabled by the firing of $t_1 \dots t_{i-1}$. Since this property only depends on discrete markings, which are preserved in ρ' , t_i is also newly enabled by the firing of $t_1 \dots t_{i-1}$ in ρ' . As a consequence, we have $\nu'_i(t_i) = 0 = \nu_i(t_i)$, thus proving that t_i is firable in ρ' , yielding a contradiction.

2. If $i \in I_{\neq 0}$. Then we have $\nu_i(t_i) \neq 0$. As already mentioned above, we have in this case that t_i is never newly enabled during $t_1 \dots t_{i-1}$ in ρ . Since the discrete markings are preserved, this conclusion holds also in ρ' . As a consequence, we have $\nu_i(t_i) = \nu(t_i) + \delta$ and $\nu'_i(t_i) = \nu(t_i) + \delta + \frac{b}{2}$, what yields $\nu'_i(t_i) = \nu(t_i) + \frac{b}{2}$. Using Property (1), we obtain $\nu'_i(t_i) \in I(t_i)$, and then t_i is fireable from (M_i, ν'_i) , contradicting the assumption.

Finally, we have obtained a contradiction in both cases, thus proving that ρ' is fireable in $\llbracket \mathcal{N} \rrbracket_s$. \square

We now consider the TPN \mathcal{N}_2 represented on Figure 5. Equipped with persistent atomic semantics, it accepts the set of timed words composed of letters a occurring before time 1. We will prove that this timed language cannot be accepted by any TPN equipped with the weak atomic semantics.




Fig. 5. The TPN \mathcal{N}_2 .

Proposition 14. *There exists no TPN \mathcal{N} (even unbounded) s.t. $\llbracket \mathcal{N} \rrbracket_A \approx \llbracket \mathcal{N}_2 \rrbracket_{PA}$.*

Proof. Assume there exists a TPN \mathcal{N} such that $\llbracket \mathcal{N} \rrbracket_A \approx \llbracket \mathcal{N}_2 \rrbracket_{PA}$. Denote by N the number of transitions of \mathcal{N} , by b the smallest positive upper bound of the intervals of \mathcal{N} , and consider a timed word $w = (a, \eta_1)(a, \eta_2) \dots (a, \eta_k)$ such that $\forall i, 1 - \frac{b}{2} < \eta_i < \eta_{i+1} < 1$, and $k \geq N + 1$.

This timed word w is recognized by $\llbracket \mathcal{N}_1 \rrbracket_{PA}$ and there exists thus a run of $\llbracket \mathcal{N} \rrbracket_A$ along w . We denote it by ρ and decompose it as follows :

$$\rho : \theta_0 \xrightarrow{d_1} \theta_1 \xrightarrow{t_a^1} \theta'_1 \xrightarrow{d_2} \theta_2 \xrightarrow{t_a^2} \theta'_2 \xrightarrow{d_3} \theta_3 \xrightarrow{t_a^3} \theta'_3 \xrightarrow{d_4} \theta_4 \xrightarrow{t_a^4} \theta'_4 \xrightarrow{d_5} \theta_5 \xrightarrow{t_a^5} \theta'_5 \xrightarrow{d_6} \theta_6 \xrightarrow{t_a^6} \theta'_6 \xrightarrow{d_7} \theta_7 \xrightarrow{t_a^7} \theta'_7 \xrightarrow{d_8} \theta_8 \xrightarrow{t_a^8} \theta'_8 \xrightarrow{d_9} \theta_9 \xrightarrow{t_a^9} \theta'_9 \xrightarrow{d_{10}} \theta_{10} \xrightarrow{t_a^{10}} \theta'_{10}$$

To obtain this decomposition we proceed as follows. We denote by t_a^i the i -th transition labelled by a . Then for each position i , we isolate the last delay step occurring before the transition t_a^i (it exists since $\eta_i > \eta_{i-1}$) and denote it by d_i . Then we gather all the internal transitions occurring between this delay step and the transition t_a^i , and denote this sequence by θ_i . The transitions between t_a^{i-1} and the delay step constitute the sequence θ'_{i-1} . In particular, the following properties hold for any position i : $\Lambda(t_a^i) = a$, $\text{Untimed}(\theta_i) = \text{Untimed}(\theta'_{i-1}) = \varepsilon$, $d_i > 0$, $\text{Duration}(\theta_i) = 0$, and t_a^i occurs at time η_i .

We claim there exists an index $i \in \{1, \dots, k\}$ such that each transition t appearing in θ_i t_a^i has already been fired since θ_0 , i.e. t also appears in $\theta_1 t_a^1 \theta'_1 \dots \theta_{i-1} t_a^{i-1} \theta'_{i-1}$. By contradiction, if it is not the case, then we can find, for each index $i \in \{1, \dots, k\}$, a transition, denoted t_i , that never appears before. The choice of k verifying $k \geq N + 1$ then implies that there exist two positions $i \neq j$ such that $t_i = t_j$, thus yielding a contradiction. We can now fix an index i verifying the above described property.

We now show that Lemma 13 can be applied to the part of ρ associated with the sequence $d_i \theta_i t_a^i$. More precisely, (M, ν) is the configuration reached after firing $\theta_0 \dots t_a^{i-1} \theta'_{i-1}$, the delay δ is equal to d_i , the sequence $t_1 \dots t_n$ corresponds to $\theta_i t_a^i$, and η is defined as $(\eta_i - d_i) - (1 - \frac{b}{2})$. In the atomic semantics, when a transition is fired, its clock is reset if it is still enabled. This property allows, together with timing constraints on the word w , to verify hypotheses (i) and (ii) of the Lemma 13. Indeed,

since each transition in $\theta_i t_a^i$ has been reset along $\theta_1 t_a^1 \theta'_1 \dots \theta'_{i-1}$, it has been reset since time η_1 . Since the global time associated with (M, ν) is equal to $\eta_i - d_i$, these valuations are bounded by above by the value $(\eta_i - d_i) - \tau_1 \leq (\eta_i - d_i) - (1 - \frac{b}{2}) = \eta$. Second, we have $\eta + \delta = \eta_i - (1 - \frac{b}{2}) < \frac{b}{2}$, as desired (this follows from the inequalities $1 - \frac{b}{2} < \eta_i < 1$).

Finally, Lemma 13 thus allows to delay of $\frac{b}{2}$ the firing of the sequence $\theta_i t_a^i$. In particular, this will produce a letter a at time $\eta_i + \frac{b}{2} > 1$. The TTS $\llbracket \mathcal{N} \rrbracket_A$ thus accepts a timed word not recognized by $\llbracket \mathcal{N}_2 \rrbracket_{PA}$, providing a contradiction. \square

Using the results of Propositions 12 and 14, we deduce that:

Theorem 15. $\mathcal{TPN}_A \subset \mathcal{TPN}_{PA}$.

4.3 About Atomic and Intermediate policies in weak and strong semantics

In this subsection, we discuss the comparison of the intermediate and atomic policies. As we will see, the situation is more complex than in the previous comparison.

On the inclusion of \mathcal{TPN}_I into \mathcal{TPN}_A . For the **strong semantics**, a construction has been proposed in [4] to transform any TPN with intermediate policy into an equivalent (w.r.t. weak timed bisimilarity) TPN with atomic semantics. A first attempt was thus to adapt this construction for the weak semantics. But studying this construction, we noticed that it is erroneous (even for the strong semantics). We present below an example exhibiting the error.




Fig. 6. A counter example to the construction of [4].

Example 2. Consider the net \mathcal{N}_3 depicted on the left of Figure 6. The application of the construction proposed in [4] leads to the net \mathcal{N}'_3 depicted on the right of Figure 6. According to [4], we should have, under the **strong semantics**, the relation $\llbracket \mathcal{N}_3 \rrbracket_I \approx \llbracket \mathcal{N}'_3 \rrbracket_A$. However, it is easy to verify that in the TTS $\llbracket \mathcal{N}'_3 \rrbracket_A$ the letter c can be read after 2 times units (with the timed word $(a, 1)(a, 1)(\tau, 1)(\tau, 1)(c, 2)$) whereas it is not possible in $\llbracket \mathcal{N}_3 \rrbracket_I$, thus proving that the construction proposed in [4] is erroneous.

This example leaves open the question of the inclusion of \mathcal{TPN}_I into \mathcal{TPN}_A for the strong semantics, and then for this semantics both inclusions are left open. For weak semantics, this inclusion is also open, but we show below that the converse inclusion is false.

Non inclusion of \mathcal{TPN}_A into \mathcal{TPN}_I . We exhibit a TPN with atomic semantics which cannot be expressed in an equivalent way by any TPN with intermediate semantics (with weak elapsing of time). This is formally stated in the Proposition below. We consider the TPN \mathcal{N}_1 represented on Figure 1. Interpreted in weak atomic semantics, the firing of the a -labelled transition does not newly enable transition labelled by c . This transition thus shares a token with transition a while preserving a time reference to the origin of global time, what is impossible in intermediate semantics.

Proposition 16. *There exists no TPN \mathcal{N} (even unbounded) such that $\llbracket \mathcal{N} \rrbracket_I \approx \llbracket \mathcal{N}_1 \rrbracket_A$.*

Proof. We proceed by contradiction and assume there exists such a TPN \mathcal{N} , and denote by N its number of transitions, and b the smallest positive upper bound of its intervals. As in the proof of Proposition 14, we first exhibit a particular execution ρ of $\llbracket \mathcal{N} \rrbracket_I$:

Lemma 17. *Let $(\eta_i)_{1 \leq i \leq k}$ be a set of timestamps such that for any $1 \leq i \leq k$, $1 - \frac{b}{2} < \eta_i < \eta_{i+1} < 1$ and $k \geq N + 1$. There exists a run ρ in $\llbracket \mathcal{N} \rrbracket_I$ of the following form:*

$$\rho : \xrightarrow{1-\frac{b}{2}}_I \xrightarrow{\theta_1}_I \xrightarrow{d_1}_I \xrightarrow{\theta'_1}_I \xrightarrow{t_a^1}_I \xrightarrow{\theta''_1}_I \cdots \xrightarrow{\theta_i}_I \xrightarrow{d_i}_I \xrightarrow{\theta'_i}_I \xrightarrow{t_a^i}_I \xrightarrow{\theta''_i}_I \cdots \xrightarrow{\theta_n}_I \xrightarrow{d_n}_I \xrightarrow{\theta'_n}_I \xrightarrow{t_a^n}_I \xrightarrow{\theta''_n}_I$$

such that for any position i , $\Lambda(t_a^i) = a$, the transition t_a^i occurs at time η_i , $d_i > 0$, $\text{Untimed}(\theta_i) = \text{Untimed}(\theta'_i) = \text{Untimed}(\theta''_i) = \varepsilon$, $\text{Duration}(\theta'_i) = \text{Duration}(\theta''_i) = 0$, and there exists a transition t_c^i , labelled by c , newly enabled by the last transition of $t_a^i \theta''_i$ and (immediately) fireable from the configuration reached after θ''_i .

Proof. We present here the details of the construction of the run ρ . This construction proceeds in three steps.

First step : construction of the structure of ρ . This step is the most involved one. We deeply use the bisimulation property between $\llbracket \mathcal{N}_1 \rrbracket_A$ and $\llbracket \mathcal{N} \rrbracket_I$ to build a first sequence. Figure 7 illustrates this construction. A dashed arrow between two configurations means that these configurations are bisimilar. The direction of this arrow indicates which implication is used to obtain the bisimilarity (from $\llbracket \mathcal{N}_1 \rrbracket_A$ to $\llbracket \mathcal{N} \rrbracket_I$, or conversely). In this figure, we omit the index I and A which should be associated to each of the step. We now detail step by step how this Figure reads. Initially, due to weak semantics, the net \mathcal{N} in intermediate semantics can choose to delay $1 - \frac{b}{2}$ time units. By bisimulation (and because there are no silent transitions in \mathcal{N}_1), the same delay leads to a bisimilar configuration in $\llbracket \mathcal{N}_1 \rrbracket_A$, thus indicated by a bottom-up dashed arrow. From this configuration one can in $\llbracket \mathcal{N}_1 \rrbracket_A$ delay $\eta_1 - (1 - \frac{b}{2})$ time units³ and then fire the a transition. By bisimulation, there exists a path in $\llbracket \mathcal{N} \rrbracket_I$, written as $u_1 t_a^1 u'_1$, leading to a bisimilar configuration and such that u_1 only contains silent transitions and is of duration exactly $\eta_1 - \eta_0$, t_a^1 is labelled by a , and u'_1 is labelled with internal actions and has a null duration. The bisimilarity is indicated by a top-down dashed arrow since it is the existence of a path in $\llbracket \mathcal{N}_1 \rrbracket_A$ that implies the existence of a path in $\llbracket \mathcal{N} \rrbracket_I$. Next, one can fire instantaneously c in $\llbracket \mathcal{N}_1 \rrbracket_A$ what implies the existence of a path $\sigma_1 t_c^1 \sigma'_1$, leading to a bisimilar configuration, and such that σ_1 and σ'_1 are labelled by internal actions and of null duration, and t_c^1 is labelled by c . Then, we use the bisimulation property in the converse direction: the configuration reached after the prefix σ_1 , which is labelled by

³ We denote by η_0 the value $1 - \frac{b}{2}$.

internal actions and of null duration, is still bisimilar with the configuration of $\llbracket \mathcal{N}_1 \rrbracket_A$ reached after the a transition, as indicated by the bottom-up dashed arrow (the only path in $\llbracket \mathcal{N}_1 \rrbracket_A$ of null duration and labelled by internal actions is the empty path). Then the same reasoning is applied from these two bisimilar configurations, and can be repeated arbitrarily many times.




Fig. 7. Using bisimulation to build ρ .

Second step : definition of θ_i, θ'_i and d_i . This step is much simpler. We simply split the paths u_i in three (eventually empty) parts. The duration of u_i is equal to $\eta_i - \eta_{i-1} > 0$, thus we can consider the last positive delay step appearing in u_i , and denote it by d_i . Then we simply write $u_i = \theta_i \cdot d_i \cdot \theta'_i$.

Third step : definition of θ''_i . In this last step, we will define the path θ''_i as a prefix of the path $u'_i \cdot \sigma_i$. This path possesses the following properties : its duration is null, it is labelled by internal actions, and there exists a c labelled transition t_c^i immediately fireable after it. The only property missing is that its last transition newly enables transition t_c^i . To obtain this property, we will show that we can find a prefix of $u'_i \cdot \sigma_i$ having this additional property. Therefore we prove the following lemma, specific to the intermediate semantics:

Lemma 18. *Let \mathcal{N} be a TPN, and consider an execution of $\llbracket \mathcal{N} \rrbracket_I$, denoted by $\rho : (M_1, v_1) \xrightarrow{t_1}_I (M_2, v_2) \xrightarrow{t_2}_I \dots (M_n, v_n) \xrightarrow{t_n}_I (M_{n+1}, v_{n+1}) \xrightarrow{t}_I$, such that for any $i \in \{1, \dots, n\}$, $\uparrow \text{enabled}_I(t, M_i, t_i) = \text{false}$.*

Then, for any $i \in \{1, \dots, n\}$, the execution $\rho_i : (M_i, v_i) \xrightarrow{t}_I \xrightarrow{t_i}_I$ exists in $\llbracket \mathcal{N} \rrbracket_I$.

Proof (of Lemma 18). By definition of the predicate $\uparrow \text{enabled}_I$, and since t is newly enabled by no t_i and fireable from (M_{n+1}, v_{n+1}) , we have that t is enabled in every M_i and, because the sequence $t_1 \dots t_n$ is instantaneous (no delay step), $v_{n+1}(t) = v_i(t)$ for any i . Thus, t is fireable from the configuration (M_i, v_i) . Moreover, we can notice, as t_i is enabled in M_i , that $\uparrow \text{enabled}_I(t, M_i, t_i) = \text{false}$ implies that t_i is still enabled after the firing of t from M_i , and has not been newly enabled: $\uparrow \text{enabled}_I(t_i, M_i, t) = \text{false}$. As a consequence, t_i is fireable from the configuration reached after firing t , thus proving the Lemma. \square

Lemma 18 entails that there exists a transition in the sequence $t_a^i u_i' \sigma_i$ which newly enables the transition t_c^i . Indeed, if it is not the case, the result of Lemma 18 implies that after the firing of u_i , one can first fire the transition t_c^i , and immediately after fire the transition t_a^i . This leads to the production of a letter a after the letter c , which is impossible in $\llbracket \mathcal{N}_1 \rrbracket_A$, thus leading to a contradiction. Finally, we define θ_i'' as the (eventually empty if t_c^i is newly enabled by t_a^i) prefix of $u_i' \cdot \sigma_i$ up to the first transition that newly enables t_c^i . This concludes the proof of the existence of the sequence ρ . \square

We now return on the proof of Proposition 16. First, we claim that there exists an index $i \in \{1, \dots, k\}$ such that each transition t appearing in $\theta_i' t_a^i \theta_i''$ has already been fired since time η_1 , i.e. t also appears in $\theta_1' t_a^1 \theta_1'' \dots \theta_{i-1}' \theta_{i-1}^{i-1} \theta_{i-1}''$. The reasoning is similar to the one of the proof of Proposition 14: by contradiction, if there exists no such position, then we can find, for each index $i \in \{1, \dots, k\}$, a transition, denoted t_i , that never appears before. The choice of k verifying $k \geq N + 1$ then implies that there exist two positions $i \neq j$ such that $t_i = t_j$, thus yielding a contradiction. We can now fix an index i verifying the above described property.

We now show that Lemma 13 can be applied to the portion of ρ associated with the sequence $\xrightarrow{d_i} \xrightarrow{\theta_i'} \xrightarrow{t_a^i} \xrightarrow{\theta_i''}$. More precisely, let (M, ν) be the configuration reached after firing $(1 - \frac{b}{2})\theta_1 \dots t_a^{i-1} \theta_{i-1}''$, the delay δ is equal to d_i , the sequence $t_1 \dots t_n$ corresponds to $\theta_i' t_a^i \theta_i''$, and η is defined as $(\eta_i - d_i) - (1 - \frac{b}{2})$. In the intermediate semantics, when a transition is fired, its clock is reset if it is still enabled. This property allows, together with timing constraints on the run ρ , to verify hypotheses (i) and (ii) of the Lemma 13. Indeed, we first have that each transition in $\theta_i' t_a^i \theta_i''$ has been reset since time τ_1 . Since the global time associated with (M, ν) is equal to $\eta_i - d_i$, these valuations are bounded by above by the value $(\eta_i - d_i) - \eta_1 \leq (\eta_i - d_i) - (1 - \frac{b}{2}) = \eta$. Second, we have $\eta + \delta = \eta_i - (1 - \frac{b}{2}) < \frac{b}{2}$, as desired (this follows from the inequalities $1 - \frac{b}{2} < \tau_i < 1$).

Lemma 13 thus allows to delay of $\frac{b}{2}$ time units the firing of thesequence $\theta_i' t_a^i \theta_i''$. Moreover, as the transition t_c^i is newly enabled by the last transition of $t_a^i \theta_i''$, and can be immediately fireable after θ_i'' in ρ , we deduce that this immediate firing can also be performed in the delayed execution. We thus obtain an execution in $\llbracket \mathcal{N} \rrbracket_{\mathcal{I}}$ with a c action following an a action after time 1, which is impossible in $\llbracket \mathcal{N}_1 \rrbracket_A$, thus yielding a contradiction. \square

5 Conclusion

We have studied in this paper the model of Time Petri Nets under a weak semantics of time elapsing, allowing any delay transition. We have first proven that for the intermediate memory policy, the set of reachable markings coincides with the reachability set of the underlying untimed Petri net. As a consequence, many verification problems are decidable for weak intermediate TPN. On the other hand, we have proven that the two other memory policies, namely atomic and persistent atomic, allow to simulate Minsky machines and thus are undecidable. Finally, we have studied expressiveness and have proven that (i) the atomic semantics is strictly less expressive than the persistent atomic one and (ii) the atomic semantics is not included in the intermediate one.

In further work, we plan to investigate properties concerning executions of weak intermediate TPN; such as time-optimal reachability, or LTL model checking. Indeed, while discrete markings are the same, the executions are different from those accepted by the underlying Petri net. Concerning expressiveness, we conjecture that intermediate and atomic semantics are uncomparable in general, and that bounded weak TPN are strictly less expressive than timed automata (without invariants).

Acknowledgments. We would like to thank Fabrice Chevalier for fruitful discussions and the anonymous reviewers for their insightful comments.

References

1. P. A. Abdulla, P. Mahata, and R. Mayr. Dense-timed Petri nets: Checking zenoness, token liveness and boundedness. *Logical Methods in Computer Science*, 3(1):1–61, 2007.
2. P. A. Abdulla and A. Nylén. Timed Petri nets and bqos. In *ICATPN'01*, volume 2075 of *LNCS*, pages 53–70. Springer, 2001.
3. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
4. B. Bérard, F. Cassez, S. Haddad, D. Lime, and O. H. Roux. Comparison of different semantics for time Petri nets. In *ATVA'05*, volume 3707 of *LNCS*, pages 293–307. Springer, 2005.
5. T. Bolognesi, F. Lucidi, and S. Trigila. From timed Petri nets to timed LOTOS. In *PSTV'90*, pages 395–408. North-Holland, 1990.
6. P. Bouyer, S. Haddad, and P.-A. Reynier. Timed Petri nets and timed automata: On the discriminating power of zeno sequences. *Information and Computation*, 206(1):73–107, 2008.
7. M. Boyer and O. H. Roux. On the compared expressiveness of arc, place and transition time Petri nets. *Fundamenta Informaticae*, 88(3):225–249, 2008.
8. D. de Frutos-Escrig, V. Valero Ruiz, and O. M. Alonso. Decidability of properties of timed-arc Petri nets. In *ICATPN'00*, volume 1825 of *LNCS*, pages 187–206. Springer, 2000.
9. S. Haar, F. Simonot-Lion, L. Kaiser, and J. Toussaint. Equivalence of timed state machines and safe time Petri nets. In *WoDES'02*, pages 119–126, 2002.
10. N. D. Jones, L. H. Landweber, and Y. E. Lien. Complexity of some problems in Petri nets. *Theoretical Computer Science*, 4(3):277–299, 1977.
11. R. M. Karp and R. E. Miller. Parallel program schemata. *Journal of Computer System Sciences*, 3(2):147–195, 1969.
12. S. R. Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *STOC'82*, pages 267–281. ACM, 1982.
13. E. W. Mayr. An algorithm for the general Petri net reachability problem. *SIAM Journal on Computing*, 13(3):441–460, 1984.
14. P. M. Merlin. *A Study of the Recoverability of Computing Systems*. PhD thesis, University of California, Irvine, CA, USA, 1974.
15. M. L. Minsky. *Computation: finite and infinite machines*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1967.
16. J. Sifakis and S. Yovine. Compositional specification of timed systems. In *STACS'96*, volume 1046 of *LNCS*, pages 347–359. Springer, 1996.
17. V. Valero Ruiz, D. de Frutos-Escrig, and F. Cuartero. On non-decidability of reachability for timed-arc Petri nets. In *PNPM'99*, pages 188–196. IEEE Computer Society Press, 1999.