

Structural translation from Time Petri Nets to Timed Automata

Franck Cassez *, Olivier H. Roux

IRCCyNICNRS UMR 6597, BP 92101, 1 rue de la Noë, 44321 Nantes Cedex 3, France

Received 27 July 2005; received in revised form 25 October 2005; accepted 13 December 2005

Available online 7 February 2006

Abstract

In this paper, we consider Time Petri Nets (TPN) where time is associated with transitions. We give a formal semantics for TPNs in terms of Timed Transition Systems. Then, we propose a translation from TPNs to Timed Automata (TA) that preserves the behavioral semantics (timed bisimilarity) of the TPNs. For the theory of TPNs this result is twofold: (i) reachability problems and more generally TCTL model-checking are decidable for bounded TPNs; (ii) allowing strict time constraints on transitions for TPNs preserves the results described in (i). The practical applications of the translation are: (i) one can specify a system using both TPNs and Timed Automata and a precise semantics is given to the composition; (ii) one can use existing tools for analyzing timed automata (like KRONOS, UPPAAL or CMC) to analyze TPNs. In this paper we describe the new feature of the tool ROMEO that implements our translation of TPNs in the UPPAAL input format. We also report on experiments carried out on various examples and compare the result of our method to state-of-the-art tool for analyzing TPNs.

© 2006 Elsevier Inc. All rights reserved.

Keywords: Time Petri Nets; Timed automata; Model-checking

1. Introduction

Petri Nets with time. The two main extensions of Petri Nets with time are Time Petri Nets (TPNs) (Merlin, 1974) and Timed Petri Nets (Ramchandani, 1974). For TPNs a transition can fire within a time interval whereas for Timed Petri Nets it fires as soon as possible. Among Timed Petri Nets, time can be considered relative to places or transitions (Sifakis, 1980; Pezzé and Young, 1999). The two corresponding subclasses namely P-Timed Petri Nets and T-Timed Petri Nets are expressively equivalent (Sifakis, 1980; Pezzé and Young, 1999). The same classes are denoted for TPNs i.e. T-TPNs and P-TPNs, but both classes of Timed Petri Nets are included in both P-TPNs and T-TPNs (Pezzé and Young, 1999). P-TPNs and T-TPNs are proved to be incomparable in (Khansa et al., 1996). Finally TPNs form a subclass of Time Stream Petri Nets

(Diaz and Senac, 1994) which were introduced to model multimedia applications.

The class T-TPNs is the most commonly-used subclass of TPNs and in this paper we focus on this subclass that will be henceforth referred to as TPN. For classical TPNs, boundedness is undecidable, and works on this model report undecidability results, or decidability under the assumption that the TPN is bounded (e.g. reachability in (Popova, 1991)). Recent work (Abdulla and Nylén, 2001; Frutos Escrig et al., 2000) consider timed arc Petri Nets where each token has a clock representing his “age”. The authors prove that coverability and boundedness are decidable for this class of timed arc Petri Nets by applying a backward exploration technique. However, they assume a lazy (non-urgent) behavior of the net: the firing of transitions may be delayed, even if that implies that some transitions are disabled because their input tokens become too old.

Verifying Time Petri Nets. The behavior of a TPN can be defined by timed firing sequences which are sequences of pairs (t, d) where t is a transition of the TPN and

* Corresponding author.

E-mail address: Franck.Cassez@irccyn.ec-nantes.fr (F. Cassez).

$d \in \mathbb{R}_{\geq 0}$. A sequence of transitions like $\omega = (t_1, d_1) (t_2, d_2) \dots (t_n, d_n) \dots$ indicates that t_1 is fired after d_1 time units, then t_2 is fired after d_2 time units, and so on, so that transition t_i is fired at absolute time $\sum_{k=1}^i d_k$. A marking M is reachable in a TPN if there is a timed firing sequence ω from the initial marking M_0 to M . Reachability analysis of TPNs relies on the construction of the so-called States Class Graph (SCG) that was introduced in (Berthomieu and Menasche, 1983) and later refined in (Berthomieu and Diaz, 1991). It has been recently improved in (Lilius, 1999) by using partial-order reduction methods.

For bounded TPNs, the SCG construction obviously solves the *marking reachability* problem (Given a marking M , “Can we reach M from M_0 ?”). If one wants to solve the *state reachability* problem (Given M and $v \in \mathbb{R}_{\geq 0}$ and a transition t , “Can we reach a marking M such that transition t has been enabled for v time units?”) the SCG is not sufficient and an alternative graph, the *strong state class graph* is introduced for this purpose in (Berthomieu and Vernadat, 2003). The two previous graphs allow for checking LTL properties. Another graph can be constructed that preserves CTL* properties. Anyway none of the previous graphs is a good¹ abstraction (accurate enough) for checking *quantitative* real-time properties e.g. “it is not possible to stay in marking M more than n time units” or “from marking M , marking M' is always reached within n time units”. The two main (efficient) tools used to verify TPNs are TINA (Berthomieu et al., 2004) and ROMEO (Gardey et al., 2005).

Timed Automata. Timed Automata (TA) were introduced by Alur and Dill (1994) and have since been extensively studied. This model is an extension of finite automata with (dense time) *clocks* and enables one to specify real-time systems.

It has been shown that model-checking for TCTL properties is decidable (Alur and Dill, 1994; Henzinger et al., 1994) for TA and some of their extensions (Bouyer et al., 2000

(Pettersson and Larsen, 2000), KRONOS (Yovine, 1997) and CMC (Laroussinie and Larsen, 1998) for model-checking TA and many real-time industrial applications have been specified and successfully verified with them.

Related work. The relationship between TPNs and TA has not been much investigated. In Sifakis and Yovine (1996) are mainly concerned with *compositionality* problems. They show that for a subclass of 1-safe Time Stream Petri Nets, the usual notion of composition used for TA is not suitable to describe this type of Petri Nets as the composition of TA. Consequently, they propose Timed Automata with Deadlines and flexible notions of compositions. In Bornot et al. (1998) the authors consider Petri Nets with deadlines (PND) that are 1-safe Petri Nets extended with clocks. A PND is a timed automaton with

deadlines (TAD) where the discrete transition structure is the corresponding marking graph. The transitions of the marking graph are subject to the same timing constraints as the transitions of the PND. The PND and the TAD have the same number of clocks. They propose a translation of safe TPN into PND with a clock for each input arc of the initial TPN. It defines (by transitivity) a translation of safe TPN into TAD (that can be considered as standard timed automata). In Cortés et al. (2000) the authors consider an extension of Time Petri Nets (PRES+) and propose a translation into hybrid automata. Correctness of the translation is not proved. Moreover the method is defined only for 1-safe nets.

In another line of work, Sava (2001) considers bounded TPN where the underlying Petri Net is not necessarily safe and proposes an algorithm to translate the TPN into a timed automaton (one clock is needed for each transition of the original TPN). However, the author does not give any proof that this translation is correct (i.e. it preserves some equivalence relation between the semantics of the original TPN and the computed TA) and neither that the algorithm terminates (even if the TPN is bounded).

Lime and Roux proposed an extension in Lime and Roux (2003) of the state class graph construction that allows to build the state class graph of a bounded TPN as a timed automaton. They prove that this timed automaton and the TPN are timed-bisimilar and they prove a relative minimality result of the number of clocks needed in the obtained automaton.

The first two approaches are structural but are limited to Petri Nets whose underlying net is 1-safe. The last two approaches rely on the computation of the entire (symbolic) state space of the TPN and are limited to bounded TPN. For example in Lime and Roux (2003), a timed automaton is indeed computed from a TPN but this requires timing information that are collected by computing a State Class Graph. Moreover, if one uses tools for analyzing the timed automaton, the timing correspondence with the TPN we started with is not easy to infer: the clocks of the timed automaton do not have a “uniform” meaning² in the TPN.

In this article, we consider a structural translation from TPN (not necessary bounded) to TA. The drawbacks of the previous translation are thus avoided because we do not need to compute a State Class Graph and have an easy correspondence of the clocks of our TA with the timing constraints on the transitions of the TPN we started with. It also extends previous results in the following directions: first we can easily prove that our translation is correct and terminates as it is a syntactic translation and it produces a timed automaton that is timed bisimilar to the

¹ The use of *observers* is of little help as it requires to specify a property as a TPN; thus it is hard to specify properties on markings.

² Assume you want to prove a property P is false on the timed automaton associated with a TPN and this happens because clock x (in the TA) has a particular value. There is no easy way of feeding this information back in terms of timing constraints on the transitions of the TPN.

TPN we started with. Notice that the timed automaton contains integer variables that correspond to the marking of the Petri Net and that it may have an unbounded number of locations. However timed bisimilarity holds even in the unbounded case. In case the Petri Net is bounded we obtain a timed automaton with a finite number of locations and we can check for TCTL properties of the original TPN. Second as it is a structural translation it does not need expensive computation (like the State Class Graph) to obtain a timed automaton. This has a practical application as it enables one to use efficient existing tools for TA to analyze the TPN.

Our contribution. We first give a formal semantics for Time Petri Nets (Merlin, 1974) in terms of Timed Transition Systems. Then we present a structural translation of a TPN into a synchronized product of timed automata that preserves the semantics (in the sense of *timed bisimilarity*) of the TPN. This yields theoretical and practical applications of this translation: (i) TCTL (Alur and Dill, 1994; Henzinger et al., 1994) model-checking is decidable for bounded TPNs and TCTL properties can now be checked (efficiently) for TPNs with existing tools for analyzing timed automata (like KRONOS, UPPAAL or CMC); (ii) allowing strict time constraints on transitions for TPNs preserves the previous result: this leads to an extension of the original TPN model for which TCTL properties can be decided; (iii) one can specify a system using both TPNs and Timed Automata and a precise semantics is given to the composition; (iv) as the translation is structural, one can use unboundedness testing methods to detect behavior leading to the unboundedness of a TPN.

Some of the above mentioned results appeared in the preliminary version of this paper (Cassez and Roux, 2004). In this extended version we have added the proofs of some theorems and more importantly a section that deals with the implementation of our approach: (i) we have implemented the structural translation of TPNs into UPPAAL TAs in the tool ROMEO; (ii) we have compared our approach to existing tools for analyzing TPNs.

Outline of the paper. Section 2 introduces the semantics of TPNs in terms of timed transition systems and the basics of TA. In Section 3 we show how to build a synchronized product of TA that is timed bisimilar to a TPN. We show how it enables us to check for real-time properties expressed in TCTL in Section 4. In Section 5 we describe the implementation of our translation in the tool ROMEO, and report on experiments we have carried out with UPPAAL to check properties of TPNs. Finally we conclude with our ongoing work and perspectives in Section 6.

2. Time Petri Nets and Timed Automata

Notations. We denote by B^A 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 pointwise extensions of their counterparts in A . For a *valuation*

$v \in A^n$, $d \in A$, $v + d$ denotes the vector $(v + d)_i = v_i + d$, and for $A' \subseteq A$, $v[A' \mapsto 0]$ denotes the valuation v' with $v'(x) = 0$ for $x \in A'$ and $v'(x) = v(x)$ otherwise. We denote $\mathcal{C}(X)$ for the *simple constraints* over a set of variables X . $\mathcal{C}(X)$ is defined to be the set of boolean combinations (with the connectives $\{\wedge, \vee, \neg\}$) of terms of the form $x - x' \bowtie c$ or $x \bowtie c$ for $x, x' \in X$ and $c \in \mathbb{N}$ and $\bowtie \in \{<, \leq, =, \geq, >\}$. Given a formula $\varphi \in \mathcal{C}(X)$ and a valuation $v \in A^n$, we denote by $\varphi(v)$ the truth value obtained by substituting each occurrence of x in φ by $v(x)$. For a transition system we write transitions as $s \xrightarrow{a} s'$ and a sequence of transitions of the form $s_0 \xrightarrow{a_1} s_1 \rightarrow \dots \xrightarrow{a_n} s_n$ as $s_0 \xRightarrow{w} s_n$ with $w = a_1 a_2 \dots a_n$.

2.1. Time Petri Nets

The model. Time Petri Nets were introduced in Merlin (1974) and extend Petri Nets with timing constraints on the firings of transitions.

Definition 1 (Time Petri Net). A Time Petri Net \mathcal{T} is a tuple $(P, T, \bullet(\cdot), (\cdot)\bullet, M_0, (\alpha, \beta))$ where: $P = \{p_1, p_2, \dots, p_m\}$ is a finite set of places and $T = \{t_1, t_2, \dots, t_n\}$ is a finite set of transitions; $\bullet(\cdot) \in (\mathbb{N}^P)^T$ 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; $\alpha \in (\mathbb{Q}_{\geq 0})^T$ and $\beta \in (\mathbb{Q}_{\geq 0} \cup \{\infty\})^T$ are respectively the earliest and latest firing time mappings.

Semantics of Time Petri Nets. The semantics of TPNs can be given in term of Timed Transition Systems (TTS) which are usual transition systems with two types of labels: discrete labels for events and positive reals labels for time elapsing.

$v \in (\mathbb{R}_{\geq 0})^n$ is a *valuation* such that each value v_i is the elapsed time since the last time transition t_i was enabled. $\mathbf{0}$ is the initial valuation with $\forall i \in [1..n], \mathbf{0}_i = 0$. A *marking* M of a TPN is a mapping in \mathbb{N}^P and if $M \in \mathbb{N}^P$, $M(p_i)$ is the number of tokens in place p_i . A transition t is *enabled* in a marking M iff $M \geq \bullet t$. The predicate $\uparrow \text{enabled}(t_k, M, t_i) \in \mathbb{B}$ is true if t_k is enabled by the firing of transition t_i from marking M , and false otherwise. This definition of enabledness is based on (Berthomieu and Diaz, 1991; Aura and Lilius, 2000) which is the most common one. In this framework, a transition t_k is *newly enabled* after firing t_i from marking M if “it is not enabled by $M - \bullet t_i$ and is enabled by $M' = M - \bullet t_i + t_i^\bullet$ ” (Berthomieu and Diaz, 1991).

Formally this gives:

$$\begin{aligned} \uparrow \text{enabled}(t_k, M, t_i) &= (M - \bullet t_i + t_i^\bullet \geq \bullet t_k) \\ &\wedge ((M - \bullet t_i < \bullet t_k) \vee (t_k = t_i)) \end{aligned} \quad (1)$$

Definition 2 (Semantics of TPN). The semantics of a TPN \mathcal{T} is a timed transition system $S_{\mathcal{T}} = (Q, q_0, \rightarrow)$ where: $Q = \mathbb{N}^P \times (\mathbb{R}_{\geq 0})^n$, $q_0 = (M_0, \mathbf{0})$, $\rightarrow \in Q \times (T \cup \mathbb{R}_{\geq 0}) \times Q$ consists of the discrete and continuous transition relations:

- the discrete transition relation is defined for all $t_i \in T$ by $(M, v) \xrightarrow{t_i} (M', v')$ iff:

$$\begin{cases} M \geq \bullet t_i \wedge M' = M - \bullet t_i + t_i^* \\ \alpha(t_i) \geq v_i \geq \beta(t_i) \\ v'_k = \begin{cases} 0 & \text{if } \uparrow \text{enabled}(t_k, M, t_i) \\ v_k & \text{otherwise} \end{cases} \end{cases}$$

- the continuous transition relation is defined for all $d \in \mathbb{R}_{\geq 0}$ by $(M, v) \xrightarrow{e(d)} (M, v')$ iff:

$$\begin{cases} v' = v + d \\ \forall k \in [1..n], (M \geq \bullet t_k \Rightarrow v'_k \leq \beta(t_k)) \end{cases}$$

A run of a time Petri Net \mathcal{T} is a (finite or infinite) path in $S_{\mathcal{T}}$ starting in q_0 . The set of runs of \mathcal{T} is denoted by $\llbracket \mathcal{T} \rrbracket$. The set of reachable markings of \mathcal{T} is denoted $Reach(\mathcal{T})$. If the set $Reach(\mathcal{T})$ is finite we say that \mathcal{T} is bounded. As a shorthand we write $(M, v) \xrightarrow{d}_e (M', v')$ for a sequence of time elapsing and discrete steps like $(M, v) \xrightarrow{e(d)} (M'', v'') \xrightarrow{e} (M', v')$.

This definition may need some comments. Our semantics is based on the common definition of (Berthomieu and Diaz, 1991; Aura and Lilius, 2000) for safe TPNs.

First, previous formal semantics (Berthomieu and Diaz, 1991; Lilius, 1999; Pezzé and Young, 1999; Aura and Lilius, 2000) for TPNs usually require the TPNs to be *safe*. Our semantics encompasses the whole class of TPNs and is fully consistent with the previous semantics when restricted to safe TPNs.³ Thus, we have given a semantics to multiple enabledness of transitions which seems the most simple and adequate. Indeed, several interpretations can be given to multiple enabledness (Berthomieu and Diaz, 1991).

Second, some variations can be found in the literature about TPNs concerning the firing of transitions. The paper (Pezzé and Young, 1999) considers two distinct semantics: Weak Time Semantics (WTS) and Strong Time Semantics (STS). According to WTS, a transition *can* be fired only in its time interval whereas in STS, a transition *must* fire within its firing interval unless disabled by the firing of others. The most commonly used semantics is STS as in Merlin (1974); Berthomieu and Diaz (1991); Pezzé and Young (1999); Aura and Lilius (2000).

Third, it is possible for the TPN to be zero or unbounded. In the case it is unbounded, the discrete component of the state space of the timed transition system is infinite. If $\forall i, \alpha(t_i) > 0$ then the TPN is non-zero and the requirement that time diverges on each run is fulfilled. Otherwise, if the TPN is bounded and at least one lower bound is 0, the zero or non-zero property can be decided (Henzinger et al., 1994) for the TPN using the equivalent timed automaton we build in Section 3.

2.2. Timed Automata and products of Timed Automata

Timed automata (Alur and Dill, 1994) are used to model systems which combine *discrete* and *continuous* evolutions.

Definition 3 (Timed Automaton). A Timed Automaton H is a tuple (N, l_0, X, A, E, Inv) where: N is a finite set of locations; $l_0 \in N$ is the initial location; X is a finite set of positive real-valued clocks; A is a finite set of actions; $E \subseteq N \times \mathcal{C}(X) \times A \times 2^X \times N$ is a finite set of edges, $e = \langle l, \gamma, a, R, l' \rangle \in E$ represents an edge from the location l to the location l' with the guard γ , the label a and the reset set $R \subseteq X$; and $Inv \in \mathcal{C}(X)^N$ assigns an invariant to any location. We restrict the invariants to conjuncts of terms of the form $c \leq r$ for $c \in X$ and $r \in \mathbb{N}$.

The semantics of a timed automaton is also a timed transition system.

Definition 4 (Semantics of a TA). The semantics of a timed automaton $H = (N, l_0, X, A, E, Inv)$ is a timed transition system $S_H = (\mathcal{Q}, q_0, \rightarrow)$ with $\mathcal{Q} = N \times (\mathbb{R}_{\leq 0})^X$, $q_0 = (l_0, \mathbf{0})$ is the initial state and \rightarrow consists of the discrete and continuous transition relations:

- the discrete transition relation is defined for all $a \in A$ by $(l, v) \xrightarrow{a} (l', v')$ iff:

$$\exists (l, \gamma, a, R, l') \in E \text{ s.t. } \begin{cases} \gamma(v) = \text{tt} \\ v' = v[R \mapsto 0] \\ Inv(l')(v') = \text{tt} \end{cases}$$

- the continuous transitions is defined for all $t \in \mathbb{R}_{\geq 0}$ by $(l, v) \xrightarrow{t} (l', v')$ iff:

$$\begin{cases} l = l' & v' = v + t \quad \text{and} \\ \forall 0 \leq t' \leq t, Inv(l)(v + t') = \text{tt} \end{cases}$$

A run of a timed automaton H is a path in S_H starting in q_0 . The set of runs of H is denoted by $\llbracket H \rrbracket$.

Product of Timed Automata. It is convenient to describe a system as a parallel composition of timed automata. To this end, we use the classical composition notion based on a *synchronization function* à la Arnold-Nivat. Let $X = \{x_1, \dots, x_n\}$ be a set of clocks, H_1, \dots, H_n be n timed automata with $H_i = (N_i, l_{i,0}, X, A, E_i, Inv_i)$. A *synchronization function* f is a partial function from $(A \cup \{\bullet\})^n \hookrightarrow A$ where \bullet is a special symbol used when an automaton is not involved in a step of the global system. Note that f is a synchronization function with renaming. We denote by $(H_1 | \dots | H_n)_f$ parallel composition of the H_i 's w.r.t. f . The configurations of $(H_1 | \dots | H_n)_f$ are pairs (\mathbf{l}, \mathbf{v}) with $\mathbf{l} = (l_1, \dots, l_n) \in N_1 \times \dots \times N_n$ and $\mathbf{v} = (v_1, \dots, v_n)$ where each v_i is the value of the clock $x_i \in X$. Then the semantics of a synchronized product of timed automata is also a timed transition system: the synchronized product can do a discrete transition if all the components agree to and time can progress in the synchronized product also if all the components agree to. This is formalized by the following definition:

³ If we accept the difference with (Lilius, 1999) in the definition of the reset instants for newly enabled transitions.




Fig. 1. Automata for the transitions and the supervisor.

Definition 5 (*Semantics of a Product of TA*). Let H_1, \dots, H_n be n timed automata with $H_i = (N_i, l_{i,0}, X, A, E_i, Inv_i)$, and f a (partial) synchronization function $(A \cup \{\bullet\})^n \hookrightarrow A$. The semantics of $(H_1 | \dots | H_n)_f$ is a timed transition system $S = (Q, q_0, \rightarrow)$ with $Q = N_1 \times \dots \times N_n \times (\mathbb{R}_{\geq 0})^X$, q_0 is the initial state $((l_{1,0}, \dots, l_{n,0}), \mathbf{0})$ and \rightarrow is defined by

- $(\mathbf{l}, \mathbf{v}) \xrightarrow{b} (\mathbf{l}', \mathbf{v}')$ if there exists $(a_1, \dots, a_n) \in (A \cup \{\bullet\})^n$ s.t. $f(a_1, \dots, a_n) = b$ and for any i we have:
 - If $a_i = \bullet$, then $\mathbf{l}'[i] = \mathbf{l}[i]$ and $\mathbf{v}'[i] = \mathbf{v}[i]$,
 - If $a_i \in A$, then $(\mathbf{l}[i], \mathbf{v}[i]) \xrightarrow{a_i} (\mathbf{l}'[i], \mathbf{v}'[i])$.
- $(\mathbf{l}, \mathbf{v}) \xrightarrow{c(i)} (\mathbf{l}, \mathbf{v}')$ if for all $i \in [1..n]$, every H_i agrees on time elapsing i.e. $(\mathbf{l}[i], \mathbf{v}[i]) \xrightarrow{c(i)} (\mathbf{l}[i], \mathbf{v}'[i])$.

We could equivalently define the product of n timed automata syntactically, building a new timed automaton from the n initial ones. In the sequel we consider a product $(H_1 | \dots | H_n)_f$ to be a timed automaton the semantics of which is timed bisimilar to the semantics of the product we have given in Definition 5.

3. From Time Petri Nets to Timed Automata

In this section, we build a synchronized product of timed automata from a TPN so that the behaviors of the two are in a one-to-one correspondence.

3.1. Translating Time Petri Nets into Timed Automata

We start with a TPN $\mathcal{T} = (P, T, \bullet(\cdot), (\cdot)\bullet, M_0, (\alpha, \beta))$ with $P = \{p_1, \dots, p_m\}$ and $T = \{t_1, \dots, t_n\}$.

Timed automaton for one transition. We define one timed automaton \mathcal{A}_i for each transition t_i , of T (see Fig. 1a). This timed automaton has one clock x_i . Also the states of the automaton \mathcal{A}_i give the state of the transition t_i : in state t the transition is enabled; in state \bar{t} it is disabled and in

Firing it is being fired. The initial state of each \mathcal{A}_i depends on the initial marking M_0 of the Petri Net we want to translate. If $M_0 \geq \bullet t_i$, then the initial state is t otherwise it is \bar{t} . This automaton updates an array of integers \mathbf{p} (s.t. $\mathbf{p}[i]$ is the number of tokens in place p_i) which is shared by all the \mathcal{A}_i 's. This is not covered by Definition 5, which is very often extended (Pettersson and Larsen, 2000) with integer variables (this does not affect the expressiveness of the model when the variables are bounded).

The supervisor. The automaton for the supervisor SU is depicted on Fig. 1b. The locations 1–3 subscripted with a “c” are assumed to be urgent or committed⁴ which means that no time can elapse while visiting them. We denote by $\Delta(\mathcal{T}) = (SU | \mathcal{A}_1 | \dots | \mathcal{A}_n)_f$ the timed automaton associated to the TPN \mathcal{T} . The supervisor's initial state is 0. Let us define the synchronization function⁵ f with $n+1$ parameters defined by

- $f(!pre, \bullet, \dots, ?pre, \bullet, \dots) = pre_i$ if $?pre$ is the $(i+1)$ th argument and all the other arguments are \bullet ,
- $f(!post, \bullet, \dots, ?post, \bullet, \dots) = post_i$ if $?post$ is the $(i+1)$ th argument and all the other arguments are \bullet ,
- $f(!update, ?update, \dots, ?update) = update$.

We will prove in the next subsection that the semantics of $\Delta(\mathcal{T})$ is closely related to the semantics of \mathcal{T} . For this we have to relate the states of \mathcal{T} to the states of $\Delta(\mathcal{T})$ and we define the following equivalence:

Definition 6 (*State Equivalence*). Let (M, v) and $((s, \mathbf{p}), \mathbf{q}, \mathbf{v})$ be, respectively, a state of $S_{\mathcal{T}}$ and a configuration.⁶ Then $(M, v) \approx ((s, \mathbf{p}), \mathbf{q}, \mathbf{v})$ if:

⁴ In SU , committed locations can be simulated by adding an extra variable: see (Tripakis, 1999) Appendix A for details.

⁵ The first element of the vector referees to the supervisor move.

⁶ $(s, \mathbf{p}) \in \{0, 1_c, 2_c, 3_c\} \times \mathbb{N}^m$ is the state of S , \mathbf{q} gives the product location of $\mathcal{A}_1 | \dots | \mathcal{A}_n$ and $\mathbf{v}[i]$, $i \in [1..n]$ gives the value of the clock x_i .

$$\begin{cases} s = 0 \\ \forall i \in [1..m], \mathbf{p}[i] = M(p_i) \\ \forall k \in [1..n], \mathbf{q}[k] = \begin{cases} t & \text{if } M \geq \bullet t_k \\ \bar{t} & \text{otherwise} \end{cases} \\ \forall k \in [1..n], \mathbf{v}[k] = v_k \end{cases}$$

3.2. Soundness of the translation

We now prove that our translation preserves the behaviors of the initial TPN in the sense that the semantics of the TPN and its translation are timed bisimilar. We assume a TPN \mathcal{T} and $S_{\mathcal{T}} = (Q, q_0, \rightarrow)$ its semantics. Let \mathcal{A}_i be the automaton associated with transition t_i of \mathcal{T} as described by Fig. 1a, SU the supervisor automaton of Fig. 1b and f the synchronization function defined previously. The semantics of $\Delta(\mathcal{T}) = (SU |_{\mathcal{A}_1} \cdots |_{\mathcal{A}_n})_f$ is the TTS $S_{\Delta(\mathcal{T})} = (Q_{\Delta(\mathcal{T})}, q_0^{\Delta(\mathcal{T})}, \rightarrow)$.

Theorem 1 (Timed Bisimilarity). *For $(M, v) \in S_{\mathcal{T}}$ and $((0, \mathbf{p}), \mathbf{q}, \mathbf{v}) \in S_{\Delta(\mathcal{T})}$ such that $(M, v) \approx ((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$ the following holds:*

$$(M, v) \xrightarrow{t_i} (M', v') \text{ iff } \begin{cases} ((0, \mathbf{p}), \mathbf{q}, \mathbf{v}) \xrightarrow{w_i} ((0, \mathbf{p}'), \mathbf{q}', \mathbf{v}') \text{ with} \\ w_i = \text{pre}_i.\text{update}.\text{post}_i.\text{update} \text{ and} \\ (M', v') \approx ((0, \mathbf{p}'), \mathbf{q}', \mathbf{v}') \end{cases} \quad (2)$$

$$(M, v) \xrightarrow{\epsilon(d)} (M', v') \text{ iff } \begin{cases} ((0, \mathbf{p}), \mathbf{q}, \mathbf{v}) \xrightarrow{\epsilon(d)} ((0, \mathbf{p}'), \mathbf{q}', \mathbf{v}') \text{ and} \\ (M', v') \approx ((0, \mathbf{p}'), \mathbf{q}', \mathbf{v}') \end{cases} \quad (3)$$

Proof. We first prove statement (2). Assume $(M, v) \approx ((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$. Then as t_i can be fired from (M, v) we have: (i) $M \geq \bullet t_i$, (ii) $\alpha(t_i) \leq v_i \leq \beta(t_i)$, (iii) $M' = M - \bullet t_i + t_i^*$, and (iv) $v'_k = 0$ if $\uparrow \text{enabled}(t_k, M, t_i)$ and $v'_k = v_k$ otherwise. From (i) and (ii) and the state equivalence we deduce that $\mathbf{q}[i] = t$ and $\alpha(t_i) \leq \mathbf{v}[i] \leq \beta(t_i)$. Hence ?pre is enabled in \mathcal{A}_i . In state 0 for the supervisor, !pre is the only possible transition. As the synchronization function f allows $(\text{!pre}, \bullet, \dots, \text{?pre}, \dots, \bullet)$ the global action pre_i is possible. After this move $\Delta(\mathcal{T})$ reaches state $((1, \mathbf{p}_1), \mathbf{q}_1, \mathbf{v}_1)$ such that for all $k \in [1..n]$, $\mathbf{q}_1[k] = \mathbf{q}[k]$, $\forall k \neq i$ and $\mathbf{q}_1[i] = \text{Firing}$. Also $\mathbf{p}_1 = \mathbf{p} - \bullet t_i$ and $\mathbf{v}_1 = \mathbf{v}$.

Now the only possible transition when the supervisor is in state 1 is an update transition where all the \mathcal{A}_i 's synchronize according to f . From $((1, \mathbf{p}_1), \mathbf{q}_1, \mathbf{v}_1)$ we reach $((2, \mathbf{p}_2), \mathbf{q}_2, \mathbf{v}_2)$ with $\mathbf{p}_2 = \mathbf{p}_1$, $\mathbf{v}_2 = \mathbf{v}_1$. For all $k \in [1..n]$, $k \neq i$, $\mathbf{q}_2[k] = t$ if $\mathbf{p}_1 \geq \bullet t_k$ and $\mathbf{q}_2[k] = \bar{t}$ otherwise. Also $\mathbf{q}_2[i] = \text{Firing}$. The next global transition must be a post_i transition leading to $((3, \mathbf{p}_3), \mathbf{q}_3, \mathbf{v}_3)$ with $\mathbf{p}_3 = \mathbf{p}_2 + t_i^*$, $\mathbf{v}_3 = \mathbf{v}_2$ and for all $k \in [1..n]$, $\mathbf{q}_3[k] = \mathbf{q}_2[k]$, $\forall k \neq i$ and $\mathbf{q}_3[i] = \bar{t}$.

From this last state only an update transition leading to $((0, \mathbf{p}_4), \mathbf{q}_4, \mathbf{v}_4)$ is allowed, with $\mathbf{p}_4 = \mathbf{p}_3$, \mathbf{v}_4 and \mathbf{q}_4 given by: for all $k \in [1..n]$, $\mathbf{q}_4[k] = t$ if $\mathbf{p}_3 \geq \bullet t_k$ and \bar{t} otherwise.

$\mathbf{v}_4[k] = 0$ if $\mathbf{q}_3[k] = \bar{t}$ and $\mathbf{q}_4[k] = t$ and $\mathbf{v}_4[k] = \mathbf{v}_1[k]$ otherwise. We then just notice that $\mathbf{q}_3[k] = \bar{t}$ iff $\mathbf{p} - \bullet t_i < \bullet t_k$ and $\mathbf{q}_4[k] = t$ iff $\mathbf{p} - \bullet t_i + t_i^* \geq \bullet t_k$. This entails that $\mathbf{v}_4[k] = 0$ iff $\uparrow \text{enabled}(t_k, \mathbf{p}, t_i)$ and with (iv) gives $v'_k = \mathbf{v}_4[k]$. As $\mathbf{p}_4 = \mathbf{p}_3 = \mathbf{p}_2 + t_i^* = \mathbf{p}_1 - \bullet t_i + t_i^* = \mathbf{p} - \bullet t_i + t_i^*$ using (iii) we have $\forall i \in [1..m]$, $M'(p_i) = \mathbf{p}_4[i]$. Hence we conclude that $((0, \mathbf{p}_4), \mathbf{q}_4, \mathbf{v}_4) \approx (M', v')$.

The converse of statement (2) is straightforward following the same steps as the previous ones.

We now focus on statement (3). According to the semantics of TPNs, a continuous transition $(M, v) \xrightarrow{\epsilon(d)} (M', v')$ is allowed iff $v = v' + d$ and $\forall k \in [1..n]$, $(M \geq \bullet t_k \Rightarrow v'_k \leq \beta(t_k))$. As $(M, v) \approx ((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$, if $M \geq \bullet t_k$ then $\mathbf{q}[k] = t$ and the continuous evolution for \mathcal{A}_k is constrained by the invariant $x_k \leq \beta(t_k)$. Otherwise $\mathbf{q}[k] = \bar{t}$ and the continuous evolution is unconstrained for \mathcal{A}_k . No constraints apply for the supervisor in state 0. Hence the result. \square

We can now state a useful corollary which enables us to do TCTL model-checking for TPNs in the next section. We write $\Delta((M, v)) = ((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$ if $(M, v) \approx ((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$, $\Delta(t_i) = \text{pre}_i.\text{update}.\text{post}_i.\text{update}$ and also $\Delta(\epsilon(d)) = \epsilon(d)$. Just notice that Δ is one-to-one and we can use Δ^{-1} as well. Then we extend Δ to transitions as: $\Delta((M, v) \xrightarrow{\epsilon} (M', v')) = \Delta((M, v)) \xrightarrow{\Delta(\epsilon)} \Delta((M', v'))$ with $\epsilon \in T \cup \mathbb{R}_{\geq 0}$ (as $\Delta(t_i)$ is a word, this transition is a four step transition in $\Delta(\mathcal{T})$). Again we can extend Δ to runs: if $\rho \in \llbracket \mathcal{T} \rrbracket$ we denote $\Delta(\rho)$ the associated run in $\llbracket \Delta(\mathcal{T}) \rrbracket$. Notice that Δ^{-1} is only defined for runs ρ of $\llbracket \Delta(\mathcal{T}) \rrbracket$, the last state of which is of the form $((0, \mathbf{p}), \mathbf{q}, \mathbf{v})$ where the supervisor is in state 0. We denote this property $\text{last}(\sigma) \models SU.0$.

Corollary 1. $(\rho \in \llbracket \mathcal{T} \rrbracket \wedge \sigma = \Delta(\rho))$ iff $(\sigma \in \llbracket \Delta(\mathcal{T}) \rrbracket \wedge \text{last}(\sigma) \models SU.0)$.

Proof. The proof is a direct consequence of Theorem 1. It suffices to notice that all the finite runs of $\Delta(\mathcal{T})$ are of the form

$$\sigma = (s_0, v_0) \xrightarrow{\delta_1} (s'_0, v'_0) \xrightarrow{w_1} (s_1, v_1) \cdots \xrightarrow{\delta_n} (s'_{n-1}, v'_{n-1}) \xrightarrow{w_n} (s_n, v_n) \quad (4)$$

with $w_i = \text{pre}_i.\text{update}.\text{post}_i.\text{update}$, $\delta_i \in \mathbb{R}_{\geq 0}$, and using Theorem 1, if $\text{last}(\sigma) \models SU.0$, there exists a corresponding run ρ in \mathcal{T} s.t. $\sigma = \Delta(\rho)$. \square

This property will be used in Section 4 when we address the problem of model-checking TCTL for TPNs.

4. TCTL model-checking for Time Petri Nets

We can now define TCTL (Henzinger et al., 1994) for TPNs. The only difference with the versions of (Henzinger et al., 1994) is that the atomic propositions usually associated to states are properties of markings. For practical applications with model-checkers we assume that the TPNs we check are bounded.

TCTL for TPNs.

$$\begin{aligned}
(M, \nu) \models \mathbf{M} \bowtie \bar{V} & \quad \text{iff} \quad M \bowtie \bar{V} \\
(M, \nu) \not\models \mathbf{false} & \\
(M, \nu) \models t_k + c \leq t_j + d & \quad \text{iff} \quad \nu_k + c \leq \nu_j + d \\
(M, \nu) \models \neg \varphi & \quad \text{iff} \quad (M, \nu) \not\models \varphi \\
(M, \nu) \models \varphi \rightarrow \psi & \quad \text{iff} \quad (M, \nu) \models \varphi \text{ implies } (M, \nu) \models \psi \\
(M, \nu) \models \varphi \exists \mathcal{U}_{\bowtie c} \psi & \quad \text{iff} \quad \exists \sigma \in \llbracket \mathcal{T} \rrbracket \text{ such that} \\
& \quad \left\{ \begin{array}{l} (s_0, \nu_0) = (M, \nu) \\ \forall i \in [1..n], \forall d \in [0, d_i], (s_i, \nu_i + d) \models \varphi \\ \left(\sum_{i=1}^n d_i \right) \bowtie c \text{ and } (s_n, \nu_n) \models \psi \end{array} \right. \\
(M, \nu) \models \varphi \forall \mathcal{U}_{\bowtie c} \psi & \quad \text{iff} \quad \forall \sigma \in \llbracket \mathcal{T} \rrbracket \text{ we have} \\
& \quad \left\{ \begin{array}{l} (s_0, \nu_0) = (M, \nu) \\ \forall i \in [1..n], \forall d \in [0, d_i], (s_i, \nu_i + d) \models \varphi \\ \left(\sum_{i=1}^n d_i \right) \bowtie c \text{ and } (s_n, \nu_n) \models \psi \end{array} \right.
\end{aligned}$$

Fig. 2. Semantics of TPN-TCTL.

Definition 7 (TCTL for TPN). Assume a TPN with n places, and m transitions $T = \{t_1, t_2, \dots, t_m\}$. The temporal logic TPN-TCTL is inductively defined by

$$\begin{aligned}
\text{TPN-TCTL} ::= M \bowtie \bar{V} | \mathbf{false} | t_k + c \leq t_j + d | \neg \varphi \\
| \varphi \rightarrow \psi | \varphi \exists \mathcal{U}_{\bowtie c} \psi | \varphi \forall \mathcal{U}_{\bowtie c} \psi \quad (5)
\end{aligned}$$

where \mathbf{M} and \mathbf{false} are keywords, $\varphi, \psi \in \text{TPN-TCTL}$, $t_k, t_j \in T$, $c, d \in \mathbb{Z}$, $\bar{V} \in (\mathbb{N} \cup \{\infty\})^n$ and⁷ $\bowtie \in \{<, \leq, =, >, \geq\}$.

Intuitively the meaning of $\mathbf{M} \bowtie \bar{V}$ is that the current marking vector is in relation \bowtie with \bar{V} . The meaning of the other operators is the usual one. We use the familiar shorthands $\mathbf{true} = \neg \mathbf{false}$, $\exists \diamond_{\bowtie c} \phi = \mathbf{true} \exists \mathcal{U}_{\bowtie c} \phi$ and $\forall \square_{\bowtie c} = \neg \exists \diamond_{\bowtie c} \neg \phi$.

The semantics of TPN-TCTL is defined on timed transition systems. Let $\mathcal{T} = (P, T, \bullet(\cdot), (\cdot)^{\bullet}, M_0, (\alpha, \beta))$ be a TPN and $S_{\mathcal{T}} = (Q, q_0, \rightarrow)$ the semantics of \mathcal{T} . Let $\sigma = (s_0, \nu_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, \nu_n) \in \llbracket \mathcal{T} \rrbracket$. The truth value of a formula φ of TPN-TCTL for a state (M, ν) is given in Fig. 2.

The TPN \mathcal{T} satisfies the formula φ of TPN-TCTL, which is denoted by $\mathcal{T} \models \varphi$, iff the first state of $S_{\mathcal{T}}$ satisfies φ , i.e. $(M_0, \mathbf{0}) \models \varphi$.

We will see that thanks to Corollary 1, model-checking TPNs amounts to model-checking timed automata.

Model-Checking for TPN-TCTL. Let us assume we have to model-check formula φ on a TPN \mathcal{T} . Our method consists in using the equivalent timed automaton $\Delta(\mathcal{T})$ defined in Section 3. For instance, suppose we want to check $\mathcal{T} \models \forall \square_{\leq 3} (\mathbf{M} \geq (1, 2))$. The check means that all the states reached within the next 3 time units will have a marking such that p_1 has more than one token and p_2 more

than 2. Actually, this is equivalent to checking $\forall \square_{\leq 3} (SU.0 \rightarrow (\mathbf{p}[1] \geq 1 \wedge \mathbf{p}[2] \geq 2))$ on the equivalent timed automaton. Notice that $\exists \diamond_{\leq 3} (\mathbf{M} \geq (1, 2))$ reduces to $\exists \diamond_{\leq 3} (SU.0 \wedge (\mathbf{p}[1] \geq 1 \wedge \mathbf{p}[2] \geq 2))$. We can then define the translation of a formula in TPN-TCTL to standard TCTL for timed automata: we denote TA-TCTL the logic TCTL for timed automata.

Definition 8 (From TPN-TCTL to TA-TCTL). Let φ be a formula of TPN-TCTL. Then the translation $\Delta(\varphi)$ of φ is inductively defined by:

$$\Delta(\mathbf{M} \bowtie \bar{V}) = \bigwedge_{i=1}^n (\mathbf{p}[i] \bowtie \bar{V}_i)$$

$$\Delta(\mathbf{false}) = \mathbf{false}$$

$$\Delta(t_k + c \bowtie t_j + d) = x_k + c \bowtie x_j + d$$

$$\Delta(\neg \varphi) = \neg \Delta(\varphi)$$

$$\Delta(\varphi \rightarrow \psi) = SU.0 \wedge (\Delta(\varphi) \rightarrow \Delta(\psi))$$

$$\Delta(\varphi \exists \mathcal{U}_{\bowtie c} \psi) = (SU.0 \rightarrow \Delta(\varphi)) \exists \mathcal{U}_{\bowtie c} (SU.0 \wedge \Delta(\psi))$$

$$\Delta(\varphi \forall \mathcal{U}_{\bowtie c} \psi) = (SU.0 \rightarrow \Delta(\varphi)) \forall \mathcal{U}_{\bowtie c} (SU.0 \wedge \Delta(\psi))$$

$SU.0$ means that the supervisor is in state 0 and the clocks x_k are the ones associated with every transition t_k in the translation scheme.

Theorem 2. Let \mathcal{T} be a TPN and $\Delta(\mathcal{T})$ the equivalent timed automaton. Let (M, ν) be a state of $S_{\mathcal{T}}$ and $((s, \mathbf{p}), \mathbf{q}, \mathbf{v}) = \Delta((M, \nu))$ the equivalent state of $S_{\Delta(\mathcal{T})}$ (i.e. $(M, \nu) \approx ((s, \mathbf{p}), \mathbf{q}, \mathbf{v})$). Then $\forall \varphi \in \text{TPN-TCTL}$:

$$(M, \nu) \models \varphi \text{ iff } ((s, \mathbf{p}), \mathbf{q}, \mathbf{v}) \models \Delta(\varphi)$$

Proof. The proof is done by structural induction on the formula of TPN-TCTL. The cases of $M \bowtie \bar{V}$, \mathbf{false} , $t_k + c \leq t_j + d$, $\neg \varphi$ and $\varphi \rightarrow \psi$ are straightforward. We give the full proof for $\varphi \exists \mathcal{U}_{\bowtie c} \psi$ (the same proof can be carried out for $\varphi \forall \mathcal{U}_{\bowtie c} \psi$).

⁷ The use of ∞ in \bar{V} allows us to handle comparisons like $M(p_1) \leq 2 \wedge M(p_2) \geq 3$ by writing $\mathbf{M} \leq (2, \infty) \wedge \mathbf{M} \geq (0, 3)$.




Fig. 3. A TPN for a Producer/Consumer example in ROMEO.

Only if part. Assume $(M, v) \models \varphi \exists \mathcal{U}_{\triangleright \infty} \psi$. Then by definition, there is a run ρ in $\llbracket \mathcal{S} \rrbracket$ s.t:

$$\rho = (s_0, v_0) \xrightarrow{d_1} (s_1, v_1) \cdots \xrightarrow{d_n} (s_n, v_n)$$

and $(s_0, v_0) = (M, v)$, $\sum_{i=1}^n d_i \bowtie c$, $\forall i \in [1..n]$, $\forall d \in [0, d_i]$, $(s_i, v_i + d) \models \varphi$ and $(s_n, v_n) \models \psi$. With **Corollary 1**, we conclude that there is a run $\sigma = \Delta(\rho)$ in $\llbracket S_{\Delta(\mathcal{S})} \rrbracket$ s.t.

$$\begin{aligned} \sigma &= ((l_0, p_0), \bar{q}_0, v_0) \Rightarrow_{w_1}^{d_1} ((l_1, p_1), \bar{q}_1, v_1) \cdots \cdots \\ &\Rightarrow_{w_n}^{d_n} ((l_n, p_n), \bar{q}_n, v_n) \end{aligned}$$

and $\forall i \in [1..n]$, $((l_i, p_i), \bar{q}_i, v_i) \approx (s_i, v_i)$ (this entails that $l_i = 0$.)

Since $(s_n, v_n) \approx ((l_n, p_n), \bar{q}_n, v_n)$, using the induction hypothesis on ψ , we can assume $(s_n, v_n) \models \psi$ iff $((l_n, p_n), \bar{q}_n, v_n) \models \Delta(\psi)$ and thus we can conclude that $((l_n, p_n), \bar{q}_n, v_n) \models \Delta(\psi)$. Moreover as $l_n = 0$ we have $((l_n, p_n), \bar{q}_n, v_n) \models SU.0 \wedge \Delta(\psi)$. It remains to prove that all intermediate states satisfy $S \ .0 \rightarrow \Delta(\varphi)$. Just notice that all the intermediate states in σ not satisfying $S \ .0$ between $((l_i, p_i), \bar{q}_i, v_i)$ and $((l_{i+1}, p_{i+1}), \bar{q}_{i+1}, v_{i+1})$ satisfy $S \ .0 \rightarrow \Delta(\psi)$.

Then we just need to prove that the intermediate states satisfying $S \ .0$, i.e. the states $((l_i, p_i), \bar{q}_i, v_i)$ satisfy $\Delta(\varphi)$. As for all $i \in [1..n]$, we have $((l_i, p_i), \bar{q}_i, v_i) \approx (s_i, v_i)$, with the induction hypothesis on φ , we have $\forall i \in [1..n]$, $((l_i, p_i), \bar{q}_i, v_i) \models \Delta(\varphi)$. Moreover, again applying **Theorem 1**, we obtain for all $d \in [0, d_i]$: $((l_i, p_i), \bar{q}_i, v_i + d) \approx (s_i, v_i + d)$; applying the induction hypothesis again we conclude that for all $d \in [0, d_i]$ $((l_i, p_i), \bar{q}_i, v_i + d) \models \Delta(\varphi)$. Hence $((l_0, p_0), \bar{q}_0, v_0) \models (SU.0 \rightarrow \varphi) \exists \mathcal{U}_{\triangleright \infty} (SU.0 \wedge \psi)$.

If part. Assume $((l_0, p_0), \bar{q}_0, v_0) \models (SU.0 \rightarrow \Delta(\varphi)) \exists \mathcal{U}_{\triangleright \infty} (SU.0 \wedge \Delta(\psi))$. Then there is a run

$$\sigma = ((l_0, p_0), \bar{q}_0, v_0) \Rightarrow_{w_1}^{d_1} ((l_1, p_1), \bar{q}_1, v_1) \cdots \cdots \Rightarrow_{w_n}^{d_n} ((l_n, p_n), \bar{q}_n, v_n)$$

with $((l_n, p_n), \bar{q}_n, v_n) \models SU.0 \wedge \Delta(\psi)$ and $\forall i \in [1..n], \forall d \in [0, d_i], ((l_i, p_i), \bar{q}_i, v_i) \models (SU.0 \rightarrow \Delta(\varphi))$. **AS** $((l_n, p_n), \bar{q}_n, v_n) \models SU.0$, we can use **Corollary 1** and we know there exists a run in $\llbracket \mathcal{S} \rrbracket$

$$\rho = \Delta^{-1}(\sigma) = (s_0, v_0) \xrightarrow{d_1} (s_1, v_1) \cdots \xrightarrow{d_n} (s_n, v_n)$$

with $\forall i \in [1..n], ((l_i, p_i), \bar{q}_i, v_i) \approx (s_i, v_i)$. The induction hypothesis on $S \ .0 \wedge \Delta(\psi)$ and $((l_n, p_n), \bar{q}_n, v_n) \models S \ .0 \wedge \Delta(\psi)$ implies $(s_n, v_n) \models \psi$. For all the intermediate states of ρ we also apply the induction hypothesis: each $((l_i, p_i), \bar{q}_i, v_i)$ is equivalent to (s_i, v_i) and all the states $(s_i, v_i + d)$, $d \in [0, d_i]$ satisfy φ . Hence $(s_0, v_0) \models \varphi \exists \mathcal{U}_{\triangleright \infty} \psi$. \square

5. Implementation

In this section we describe some properties of our translation and important implementation details. Then we report on examples we have checked using our approach and UPPAAL.

5.1. Translation to UPPAAL input format

The first step in using our approach is to translate an existing TPN into a product of TA. For this we use the tool ROMEO (Gardey et al., 2005) that has been developed for the analysis of TPNs (state space computation and “on the fly” model-checking of reachability properties with a zone-based forward method and with the State Class Graph method). ROMEO has a GUI (see Fig. 3) to “draw” Timed Petri Nets and now an *export to UPPAAL* feature that implements our translation of a TPN into the equivalent TA in UPPAAL input format.⁸

⁸ The last version of UPPAAL (3.4.7) is required to read the files produced by ROMEO.




Fig. 4. A UPPAAL template (from “Export to UPPAAL” feature of ROMEO).

The textual input format for TPN in ROMEO is XML and the timed automaton is given in the “.xta” UPPAAL input format.⁹ The translation gives one timed automaton for each transition and one automaton for the supervisor *SU* as described in Section 3. The automata for each transition update an array of integers $M[i]$ (which is the number of tokens¹⁰ in place i in the original TPN). For example, the enabledness and firing conditions of a transition t_i , such that $\bullet t_i = (1, 0, 0)$ and $t_i \bullet = (0, 0, 1)$, are respectively implemented by $M[0] \geq 1$ and $M[2] := M[2] + 1$. Instead of generating one *template automaton* for each transition, we generate as many templates as types of transitions in the original TPN: the type of a transition is the number of input places and output places. For the example of Fig. 3, there are only three types of transitions (one input place to one output place, one to two and two to one) and three templates in the UPPAAL translation. Then one of these templates is instantiated for each transition of the TPN we started with. An example of a UPPAAL template for transitions having one input place and one output place is given in Fig. 4; integers $B1$ and $F1$ give respectively the index of the unique input place of the transition, and the index of the output place. The timing constraints of the transition are given by $dmin$ and $dmax$. We can handle as well transitions with input and output arcs with arbitrary weights (on the examples of Fig. 4 the input and output weights are 1).

In our translation, each transition of the TPN is implemented by a TA with one clock. The synchronized product thus contains as many clocks as the number of transitions of the TPN. At first sight one can think that the translation we have proposed is far too expensive w.r.t. to the number of clocks to be of any use when using a model-checker like UPPAAL: indeed the model-checking of TA is exponential in

the number of clocks. Nevertheless we do not need to keep track of all the clocks as many of them are not useful in many states.

5.2. Inactive clocks

When a transition in a TPN is disabled, there is no need to store the value of the clock for this transition: this was already used in the seminal paper of Berthomieu and Diaz (1991). Accordingly when the TA of a transition is in location \bar{t} (not enabled) we do not need to store the value of the clock: this means that many of the clocks can often be disregarded.

In UPPAAL there is a corresponding notion of *inactive clock*:

Definition 9 (*ppaal inactive clock*). Let \mathcal{A} be a timed automaton. Let x be a clock of \mathcal{A} and ℓ be a location of \mathcal{A} . If on all path starting from (ℓ, v) in $S_{\mathcal{A}}$, the clock x is always reset before being tested then the clock x is inactive in location ℓ . A clock is active if it is not inactive.

A consequence of the notion of inactive clocks in UPPAAL is that at location ℓ the DBM that represents the constraints on the clocks will only contain the active clocks. The next proposition (which is easy to prove on the timed automaton of a transition) states that our translation is *effective* w.r.t. active clocks reduction *i.e.* that when a TA of a transition is not in state t (enabled) the corresponding clock is considered inactive by UPPAAL.

Proposition 1. Let \mathcal{A}_i be the timed automaton associated with transition t_i , of a TPN \mathcal{T} (see Fig. 1). The clock x_i of \mathcal{A}_i is inactive in locations *Firing* and \bar{t} .

The current version of UPPAAL (3.4.7) computes active clocks syntactically for each automaton of the product. When the product automaton is computed “on the fly” (for verification purposes), the set of active clocks for a product location is simply the union of the set of active clocks of each component. Again without difficulty we obtain the following theorem:

Theorem 3. Let \mathcal{T} be a TPN and $\Delta(\mathcal{T})$ the equivalent product of timed automata (see Section 3). Let M be a (reachable) marking of $S_{\mathcal{T}}$ and ℓ the equivalent¹¹ location in $S_{\Delta(\mathcal{T})}$. The number of active clocks in ℓ is equal to the number of enabled transitions in the marking M .

Thanks to this theorem and to the *active clocks reduction* feature of UPPAAL the model-checking of TCTL properties on the network of timed automata given by our translation can be efficient. Of course there are still examples with a huge number of transitions, all enabled at any time that we will not be able to analyze, but those examples cannot be handled by any existing tool for TPN.

⁹ See www.uppaal.com for further information about UPPAAL.

¹⁰ The actual meaning of $M[i]$ is given by a table that is available in the ROMEO tool via the “Translate/Indices \Rightarrow Place/Transition” menu; the table gives the name of the place represented by $M[i]$ as well as the corresponding information for transitions.

¹¹ See Definition 6.

In the next subsection we apply our translation to some recent and non trivial examples of TPNs that can be found in Gardey et al. (2005).

5.3. Tools for analyzing TPNs

One feature of ROMEO is now to export a TPN to UPPAAL or KRONOS but it was originally developed to analyze directly TPNs and has many built-in capabilities: we refer to ROMEO STD for the tool ROMEO with these capabilities (Gardey et al., 2005). TINA (Berthomieu et al., 2004) is another state-of-the-art tool to analyze TPNs with some more capabilities than ROMEO STD: it allows to produce an Atomic State Class Graph (ASCG) on which CTL* properties can be checked. Using ROMEO STD or TINA is a matter of taste as both tools give similar results on TPNs.

Table 1 gives a comparison in terms of the classes of property (LTL, CTL, TCTL, Liveness) the tools can handle. The columns UPPAAL and KRONOS in ROMEO give the combined capabilities obtained when using our structural translation and the corresponding (timed) model-checker.

Regarding time performance ROMEO STD and TINA give almost the same results. Moreover with ROMEO STD and TINA, model-checking LTL or CTL properties will usually be faster than using ROMEO + UPPAAL: those tools implement efficient algorithms to produce the (A)SCG needed to perform LTL or CTL model-checking. On one hand it is to be noticed that both ROMEO STD and TINA need (1) to produce a file containing the (A)SCG; and then (2) to run a model-checker on the obtained graph to check for the (LTL, CTL or CTL*) property. This can be prohibitive on very large examples as the ones we use in Table 2.

Table 1
What can we do with the different tools and approaches?

	TINA	ROMEO		
		ROMEO STD	ROMEO (TPN to TA)	
			UPPAAL	KRONOS
Marking reachability	Compute marking graph	On the fly computation (zone-based method)	UPPAAL-TCTL ^c	TCTL
LTL	SCG ^a + MC ^b	SCG ^a + MC ^b		
CTL	(CTL*) ASCG ^a + MC ^b	–		
Q-liveness ^c	–		UPPAAL-Liveness ^d	
TCTL			UPPAAL-TCTL ^d	

^a SCG = Computation of the State Class Graph; ASCG = of the atomic SCG.

^b MC = requires the use of a Model-Checker on the SCG.

^c Include response properties like $\forall \square(\varphi \Rightarrow \forall \diamond \Psi)$ where φ or Ψ can contain clock constraints.

^d UPPAAL implements a subset of TCTL and a special type of liveness defined by formulas of the form $\forall \square(\varphi \Rightarrow \forall \diamond \Psi)$.

Table 2
Experimental results: checking a safety property (that is true)

	File	#places	#trans.	ROMEO + UPPAAL ^a		TINA ^b	
				Time (s)	Memory (MB)	Time (s)	Memory (MB)
Real time classical examples	oex1	12	12	0.04	1.3	0.02	1.3
	oex5 ^c	29	23	0.32	3.7	0.00	1.3
	oex7	22	20	1.0	8.9	3.66	162.5
	oex8	31	21	9.50	10	3.53	160
	t001	44	39	5.92	10.5	0.13	13.3
	t3 ^c	56	50	1.18	5	0.0	1.3
Producers consumers	P3C5	14	13	0.47	5	0.13	7
	P4C5	15	14	1.41	7.5	0.27	20
	P5C5	16	15	1.22	7.5	0.24	16
	P6C7	21	20	26.33	50.8	3.8	111
	P10C10	32	31	1.86	15	0.60	18
Large examples	Gros0	80	72	416	74	20	458
	Gros1	89	80	780	52	N/A ^c	>1300 ^d
	Gros2	116	105	5434	88	N/A ^c	>1300 ^d
	Gros3	143	130	36300	150	N/A ^c	>1300 ^d


^a We run `verifyta -q -H 65536,65536 -S 1 -s file "property P"` which is the default.

^b We run `tina -W -s 0 file` which is the default and generates the SCG.

^c As the memory consumption is very high, swapping becomes dominant and after 300 s when the 1 GB RAM is full; the process was killed after 36,000 s and within this time only 800 s were used by TINA and the rest was spent on I/O by the swap process.

^d This is virtual memory usage when the process was killed.

^e In this case UPPAAL uses more memory than TINA: this may be due to memory allocation at the beginning of the process.

Fig. 5. The TPN \mathcal{T}_g .

On the other hand neither ROMEO STD nor TINA are able to check quantitative properties such as quantitative liveness (like property of Eq. (6) below) and TCTL which in general cannot be encoded with an observer (when this possible we can translate such a quantitative property into a problem of marking reachability).

5.4. Checking a liveness quantitative property on \mathcal{T}_g

Let us consider the TPN \mathcal{T}_g of Fig. 5. The response (liveness) property,

$$\begin{aligned} \forall \square (M[1] > 0 \wedge M[3] > 0 \wedge T_{1.x} > 3) \\ \Rightarrow \forall \diamond (M[2] > 0 \wedge M[4] > 0) \end{aligned} \quad (6)$$

where $M[i]$ is the marking of the place P_i , cannot be checked with TINA (nor with ROMEO STD) and can easily be checked with our method using the translation and UPPAAL. This property means that if we do not fire T_1 before 3 t.u. then it is unavoidable that at some point in the future there is a marking with a token in P_2 and in P_4 . In UPPAAL we can use the response property template $P \rightarrow Q$ which corresponds to $\forall \square (P \Rightarrow \forall \diamond Q)$. Using our TPN-TCTL translation we obtain:

$$\begin{aligned} (SU.O \text{ and } M[1] > 0 \text{ and } M[3] > 0 \text{ and } T_{1.x} > 3) \\ \rightarrow (SU.O \text{ and } M[2] > 0 \text{ and } M[4] > 0) \end{aligned}$$

To our knowledge, the translation of TPNs to TA implemented in ROMEO is currently the only existing method allowing the checking of these properties (quantitative liveness and the UPPAAL subset of TCTL) on TPNs.

5.5. Experimental results

As a preamble we just point out that our translation is syntactic and the time to translate a TPN into an equivalent product of TA is negligible. This is in contrast with the method used in TINA and ROMEO STD where the whole state space has to be computed in order to build some graph (usually very large) and later on, a model-checker has to be used to check the property on the graph. In Table 2 the time column for TINA refers to the time needed to generate the SCG whereas in the time column for UPPAAL we

give the overall time to check a property. The tests have been performed on a Pentium-PC 2GHz with 1 GB of RAM running LINUX and with TINA 2.7.2 and UPPAAL 3.4.7 and ROMEO 2.5.0. Hereafter we give the results on three types of examples: classical cyclic or periodic synchronized tasks, producers/consumers and large real-time examples. For all these examples we check a safety property that is true so that UPPAAL must explore the whole state space of the TPN; we can then compare the computation time and memory requirements for UPPAAL and TINA.

5.6. Cyclic and periodic tasks

This set of examples we call *Real-time classical* examples consists of cyclic and periodic synchronized tasks (like the alternate bit protocol modeled in Berthomieu and Diaz (1991)). On these examples we check the *k-boundedness* property of the TPN and also the property P : $\forall \square (SU.O \Rightarrow M[1] \leq 10)$. *k-boundedness* is checked by setting the maximum value¹² of the array M to k : in case the bound is hit for one transition, the module `verifyta` of UPPAAL will issue a warning “State discarded” indicating which $M[i]$ was above the bound k . If no warning is issued, all the places are bounded by k . We check property P to force UPPAAL to explore the whole state space of the system. In the sequel we will check for 20-boundedness of the TPNs. If the result is “Property P is satisfied” and no warning “State discarded” has been issued we are sure that (1) the TPN is bounded and (2) UPPAAL has scanned the whole state space of the TPN. The results of the use of UPPAAL to check these properties are given in Table 2. All the examples of this category are bounded and satisfy property P .

Except two cases (`oex5`, `t3`) UPPAAL uses less memory than TINA whereas TINA is faster. This can be due to initial allocation of memory. Also as those case studies are rather small the measured time may not be meaningful. As we will see in the section *Large examples* below, it turns out that memory is the limiting factor when we want to model-check large examples.

5.7. Producers/consumers examples

We also experimented on *i-Producers/j-Consumers* examples (`PiCj`) where the degree of concurrency is usually very high. The results are given in Table 2. The time required to check a property with UPPAAL depends on the number of clocks simultaneously enabled in the TPN. Thus the producers-consumers `P6C7` for which there are always at least 10 transitions simultaneously enabled is a particularly unfavourable case. Of course, it is difficult to generalize from these results since they are highly dependent on the TPN and on the property that we check on it. However, they give an idea of the time and space required for the ver-

¹² Thus we check that for all place i , $M[i] \leq k$.

ification and also shows that our method can be used on non trivial examples. All these examples are bounded and P is satisfied. Again UPPAAL uses far less memory than TINA.

5.8. Large examples

These examples are very large and with TINA we can only compute the SCG for the first one. For the larger ones, TINA uses too much memory and after a while swapping becomes predominant. For the example `Gros1`, after 10 h the TINA process had only used 300 s of CPU time and was using almost all the virtual memory. No output was produced. The same happened for the examples `Gros2`, `Gros3`. On these examples, UPPAAL was able to check boundedness and the property P even if the last example (`Gros3`) required ten hours of computation: the maximum amount of memory needed is 150 MB and no swapping occurred. As it is generally admitted in the verification community, memory is the limiting factor because you cannot afford many Gigabytes of memory whereas you can wait for a couple of hours. In this way our method seems to be helpful in many cases and worth using.

6. Conclusion

In this paper, we have presented a structural translation from TPNs to TA. Any TPN \mathcal{T} and its associated TA $\Delta(\mathcal{T})$ are timed bisimilar.

Such a translation has many theoretical implications. Most of the positive theoretical results on TA carry over to TPNs. The class of TPNs can be extended by allowing strict constraints (open, half-open or closed intervals) to specify the firing dates of the transitions; for this extended class, the following results follow from our translation and from [Theorem 1](#):

- TCTL model checking is decidable for bounded TPNs. Moreover efficient algorithms used in UPPAAL ([Pettersson and Larsen, 2000](#) and [KRONOS \(Yovine, 1997\)](#)) are exact for the class of TA obtained with our translation (see recent results by [Bouyer \(2003\)](#));
- it is decidable whether a TA is non-zero or not ([Henzinger et al., 1994](#)) and thus our result provides a way to decide non-zeroness for bounded TPNs;
- lastly, as our translation is structural, it is possible to use a model-checker to find sufficient conditions of unboundedness of the TPN.

These results enable us to use algorithms and tools developed for TA to check quantitative properties on TPNs. For instance, it is possible to check real-time properties expressed in TCTL on bounded TPNs. The tool [ROMEO \(Gardey et al., 2005\)](#) that has been developed for the analysis of TPN (state space computation and “on the fly” model-checking of reachability properties) imple-

ments our translation of a TPN into the equivalent TA in UPPAAL input format.

Our approach turns out to be a good alternative to existing methods for verifying TPNs:

- with our translation and UPPAAL we were able to check safety properties on very large TPNs that cannot be handled by other existing tools;
- we also extend the class of properties that can be checked on TPNs to real-time quantitative properties.

Note also that using our translation, we can take advantage of all the features of a tool like UPPAAL: looking for counter examples is usually much faster than checking a safety property. Moreover if a safety property is false, we will obtain a counter example even for unbounded TPNs (if we use breadth-first search).

References

- Abdulla, P.A., Nylén, A., 2001. Timed Petri nets and BQOs. In: ICATPN'01/LNCS, vol. 2075. Springer-Verlag, pp. 53–72.
- Alur, R., Dill, D., 1994. A theory of timed automata. *Theoretical Computer Science B* 126, 183–235.
- Aura, T., Lilius, J., 2000. A causal semantics for time Petri nets. *Theoretical Computer Science* 243 (1–2), 409–447.
- Berthomieu, B., Diaz, M., 1991. Modeling and verification of time dependent systems using time Petri nets. *IEEE Transactions on Software Engineering* 17 (3), 259–273.
- Berthomieu, B., Menasche, M., 1983. An enumerative approach for analyzing time Petri nets. In: *Information Processing IFIP Congress Series*, vol. 9. Elsevier Science Publishers, pp. 41–46.
- Berthomieu, B., Vernadat, F., 2003. State class constructions for branching analysis of time Petri nets. In: TAGAS'2003, LNCS, vol. 2619, pp. 442–457.
- Berthomieu, B., Ribet, P.-O., Vernadat, F., 2004. The tool TINA—construction of abstract state spaces for Petri nets and time Petri nets. *International Journal of Production Research* 4 (12).
- Bornot, S., Sifakis, J., Tripakis, S., 1998. Modeling urgency in timed systems. In: COMPOS'97, Malente, Germany, LNCS, vol. 1536, pp. 103–129.
- Bouyer, P., 2003. Untameable timed automata! In: STAGS'2003, LNCS, vol. 2607, pp. 620–631.
- Bouyer, P., Dufourd, C., Fleury, E., Petit, A., 2000. Are timed automata updatable? In: CAV'2000, LNCS, vol. 1855, pp. 464–479.
- Cassez, F., Roux, O.H., 2004. Structural translation of time Petri nets into timed automata. In: Huth, M. (Ed.), *Workshop on Automated Verification of Critical Systems (AVoCS'04)*, *Electronic Notes in Computer Science*. Elsevier.
- Cortés, L.A., Eles, P., Peng, Z., 2000. Verification of embedded systems using a Petri net based representation. In: *13th International Symposium on System Synthesis (ISSS'2000)*, Madrid, Spain, pp. 149–155.
- Diaz, M., Senac, P., 1994. Time stream Petri nets: a model for timed multimedia information. In: ATPN'94, LNCS, vol. 815, pp. 219–238.
- de Frutos Escrig, D., Valero Ruiz, V., Marroquín Alonso, O., 2000. Decidability of properties of timed-arc Petri nets. In: ICATPN'00, Aarhus, Denmark, LNCS, vol. 1825, pp. 187–206.
- Gardey, G., Lime, D., Magnin, M., Roux, O.H., 2005. Roméo: a tool for analyzing time Petri nets. In: *17th International Conference on Computer Aided Verification (CAV'05)*, Edinburgh, Scotland, UK. *Lecture Notes in Computer Science*, vol. 3576, Springer. The tool can be freely downloaded from <http://romeo.rts-software.org>.

- Henzinger, T.A., Nicollin, X., Sifakis, J., Yovine, S., 1994. Symbolic model checking for real-time systems. *Information and Computation* 111 (2), 193–244.
- Khansa, W., Denat, J.P., Collart-Dutilleul, S., 1996. P-time Petri nets for manufacturing systems. In: WODES'96, Scotland, pp. 94–102.
- Laroussinie, F., Larsen, K.G., 1998. CMC: a tool for compositional model-checking of real-time systems. In: FORTE-PSTV'98. Kluwer Academic Publishers, pp. 439–456.
- Lilius, J., 1999. Efficient state space search for time Petri nets. *Electronic Notes in Theoretical Computer Science* 18.
- Lime, D., Roux, O.H., 2003. State class timed automaton of a time Petri net. In: PNPM'03. IEEE Computer Society.
- Merlin, P.M., 1974. A study of the recoverability of computing systems. Ph.D. thesis, University of California, Irvine, CA.
- Pettersson, P., Larsen, K.G., 2000. UPPAAL2k. *Bulletin of the European Association for Theoretical Computer Science* 70, 40–44.
- Pezz , M., Young, M., 1999. time Petri nets: a primer introduction. Tutorial presented at the Multi-Workshop on Formal Methods in Performance Evaluation and Applications, Spain, September.
- Popova, L., 1991. On time Petri nets. *Journal Information Processing and Cybernetics*, EIK 27 (4), 227–244.
- Ramchandani, C., 1974. Analysis of asynchronous concurrent systems by timed Petri nets. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Sava, A.T., 2001. Sur la synth se de la commande des syst mes    v nements discrets temporis s. Ph.D. thesis, INPG, Grenoble, France, November.
- Sifakis, J., 1980. Performance evaluation of systems using nets. In: *Net Theory and Applications, Advanced Course on General Net Theory of Processes and Systems*, Hamburg, LNCS, vol. 84, pp. 307–319.
- Sifakis, J., Yovine, S., 1996. Compositional specification of timed systems. In: STACS'96, LNCS, vol. 1046, pp. 347–359.
- Tripakis, S., 1999. Timed diagnostics for reachability properties. In: Cleaveland, W.R. (Ed.), (TACAS'99), LNCS, vol. 1579, pp. 59–73.
- Yovine, S., 1997. Kronos: a verification tool for real-time systems. *Journal of Software Tools for Technology Transfer* 1 (1/2), 123–133.