

Decidable Classes of Unbounded Petri Nets with Time and Urgency

Sundararaman Akshay, Blaise Genest, Loïc Hélouët

► **To cite this version:**

Sundararaman Akshay, Blaise Genest, Loïc Hélouët. Decidable Classes of Unbounded Petri Nets with Time and Urgency. Fabrice Kordon and Daniel Moldt. Application and Theory of Petri Nets and Concurrency - 37th International Conference, PETRI NETS, 2016, Jun 2016, Torun, Poland. Springer, Lecture Notes in Computer Science, pp.301 - 322, 2016, Application and Theory of Petri Nets and Concurrency - 37th International Conference, PETRI NETS, 2016. <<http://acsd2016.mat.umk.pl/>>. <10.1007/978-3-319-39086-4_18>. <hal-01379414>

HAL Id: hal-01379414

<https://hal.inria.fr/hal-01379414>

Submitted on 11 Oct 2016

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.

Decidable classes of unbounded Petri nets with time and urgency

S. Akshay¹, B. Genest², and L. Hélouët³

¹ Department of Computer Science and Engineering, IIT Bombay, India.

akshayss@cse.iitb.ac.in

² CNRS, IRISA, Rennes, France

blaise.genest@irisa.fr

³ INRIA, Campus de Beaulieu, Rennes, France

loic.helouet@inria.fr

Abstract. Adding real time information to Petri net models often leads to undecidability of classical verification problems such as reachability and boundedness. For instance, models such as Timed-Transition Petri nets (TPNs) [22] are intractable except in a bounded setting. On the other hand, the model of Timed-Arc Petri nets [26] enjoys decidability results for boundedness and control-state reachability problems at the cost of disallowing urgency (the ability to enforce actions within a time delay). Our goal is to investigate decidable classes of Petri nets with time that capture some urgency and still allow unbounded behaviors, which go beyond finite state systems.

We present, up to our knowledge, the first decidability results on reachability and boundedness for Petri net variants that combine unbounded places, time, and urgency. For this, we introduce the class of Timed-Arc Petri nets with restricted Urgency, where urgency can be used only on transitions consuming tokens from bounded places. We show that control-state reachability and boundedness are decidable for this new class, by extending results from Timed-Arc Petri nets (without urgency) [2]. Our main result concerns (marking) reachability, which is undecidable for both TPNs (because of unrestricted urgency) [20] and Timed-Arc Petri Nets (because of infinite number of “clocks”) [25]. We obtain decidability of reachability for unbounded TPNs with restricted urgency under a new, yet natural, timed-arc semantics presenting them as Timed-Arc Petri Nets with restricted urgency. Decidability of reachability under the intermediate marking semantics is also obtained for a restricted subclass.

1 Introduction

Petri nets are a simple yet powerful formalism modeling distributed systems. Several extensions have been proposed to enrich them with timing constraints, and allow specification of real-time behaviors. We first discuss the decidability and expressivity of two main variants: *Timed-Transition Petri Nets (TPNs)* [22] and *Timed-Arc Petri Nets* [26].

TPNs can constrain each transition with a timing interval. To be fireable, a transition needs to have been enabled for an amount of time in the given interval [22]. Further, when a transition has been enabled for the maximal amount of time according to its

associated interval, it must fire. This is called *urgency*. Formally, a (continuous, positive valued) *clock* is associated to each transition. Hence the number of such clocks is bounded by the number of transitions. Although the number of clocks is bounded, most problems (reachability, control-state reachability, boundedness) are undecidable for TPNs [20], as two counter machines can easily be encoded. To obtain decidability, usually one has to either restrict to bounded TPNs [7], where the number of tokens in any place is bounded, or give up urgency [24]. In the latter case, the untimed language of a TPN without urgency, also known as its weak-time semantics, is the language of the associated Petri net without timing constraints, weakening the interest of TPNs.

Timed-Arc Petri Nets, also called Timed Petri Nets, associate a (continuous, positive valued) age to each token [26, 2]. The number of continuous values is thus a priori unbounded. Each arc from a place to a transition can be constrained by a timing interval, meaning that only tokens with age in the interval can be consumed by this transition. Timed-Arc Petri Nets as explained in [18, 2] cannot encode urgency. Although the number of token ages is unbounded, the theory of well structured transition systems [17] can be applied because of monotonicity (a token is allowed to stay forever at a place). Thus, control-state reachability (whether a place can be filled with at least one token) and boundedness (whether the number of tokens in places are always bounded) are decidable for Timed-Arc Petri Nets [2]. However, the (marking) reachability problem (whether a particular marking is ever reachable) is undecidable [25].

The two models have incomparable expressive power. TPNs can produce a token exactly every unit of time using urgency, while Timed-Arc Petri Nets cannot. On the other hand, Timed-Arc Petri Nets can express *latency* requirements, while TPNs cannot: indeed, TPNs (under the intermediate marking semantics) cannot track [9, 6] the ages of an unbounded number of tokens (having slightly different ages) and consume each of them with a delay or latency of at least two time units after their creation.

Our goal in this paper is to examine the trade-off between expressivity and decidability in this setting of unbounded Petri nets with time. We start by considering a framework which is expressive enough to specify both these characteristics of *latency* and *urgency*. We aim to identify subclasses which are decidable while retaining at least a restricted form of this expressivity. To do this, we introduce *Timed-Arc Petri Nets with Urgency*, extending Timed-Arc Petri Nets with explicit urgency requirements, à la Merlin [22], forcing transitions to fire if they remains enabled for long enough.

Unsurprisingly, most problems are undecidable as soon as urgency is used on unbounded places (Proposition 1, and [19]). In earlier works, decidability results have been obtained by either imposing a bound on the number of tokens (e.g., [15, 7]) or removing urgency completely (e.g., [21, 24]). Here, we consider classes of Timed-Arc Petri Nets and of TPNs *with restricted Urgency* to obtain decidability. More specifically, transitions consuming tokens exclusively from bounded places can use urgency; other transitions consuming tokens from at least one unbounded place do not have urgency constraints. Using restricted urgency does *not* make the untimed language of a TPN with restricted Urgency the same as the language of the associated untimed Petri Net. Thus, these classes with restricted Urgency differ from TPNs with weak-time semantics [24], where all urgency constraints are ignored.

We present to our knowledge the first decidability results for a Petri net variant combining time, urgency and unbounded places. First, for the general class of Timed-Arc Petri Nets with restricted Urgency, we obtain decidability of *control-state reachability* (Theorem 1), i.e., whether a given place can ever be filled, and of boundedness. This extends decidability results [2] on Timed-Arc Petri Nets (without urgency). Our main result concerns the *decidability of (marking) reachability*. Reachability is undecidable for Timed-Arc Petri Nets (without any urgency), due to the presence of unboundedly many “clocks” (timed tokens) [25], and also for TPNs (because of unrestricted urgency) [20]. This leads us to consider TPNs *with restricted urgency*, which inherently use a bounded number of “clocks”. We define a new *timed-arc semantics* for TPNs, presenting them as a subclass of Timed-Arc Petri Nets with Urgency, in the spirit of the time on token semantics of [11]. We then obtain our *main* result: reachability is decidable for TPNs with restricted Urgency under our new timed-arc semantics (Theorem 2). This allows us to decide reachability for channel systems with specified latency assuming that there is a bound on the *throughput* of the channel (i.e., on the number of messages transferred per unit of time). While our proof for deciding reachability does not adapt to the intermediate marking semantics, we obtain decidability of reachability under the intermediate marking semantics for the subclass of TPNs with restricted constraints (Theorem 3). This class forbids specifying upper and lower bounds on transitions leaving unbounded places. We summarize the decidability and expressivity results in the table below.

| Class of (unbounded) systems | Decidability | | Expressivity | |
|---|--------------|---------------------|--------------|----------|
| | Reachability | Control-state Reach | Urgency | Latency |
| <i>Timed-Arc Petri Nets with Urgency</i> | × | × | ✓ | ✓ |
| <i>Timed-Arc Petri Nets with restricted Urgency</i> | × | ✓ | <i>R</i> | ✓ |
| Timed-Arc Petri Nets | × | ✓ | × | ✓ |
| TPNs | × | × | ✓ | × |
| <i>TPNs with restricted constraints</i> | ✓ | ✓ | <i>R</i> | × |
| <i>TPNs with restricted Urgency under new timed-arc semantics</i> | ✓ | ✓ | <i>R</i> | <i>R</i> |

Table 1. Classes of systems and their associated decidability and expressivity. The italicized rows are new results in this paper. *R* stands for restricted form of expressivity.

Related work. In [19], Timed-Arc Petri Nets were extended with urgent transitions and place invariants. In contrast to our model where a timed or discrete move is always allowed in any configuration, deadlocked configurations can be reached in [19], where no discrete move is possible, and elapsing time is forbidden. Further, urgent transitions of [19] must fire as soon as they are enabled, which corresponds to the special case of having urgency 0 in our model. Urgency has also been modeled using Black transitions in generalized stochastic nets [4] and priorities in [8], but these nets cannot model latency constraints. For TPNs, the alternative multiple server semantics [9] has been proposed to model latency, but this makes the number of clocks unbounded.

Our focus in this paper is to address decidability issues. We obtain decidability for systems with (restricted) urgency *and* unbounded places. As far as we know, in all earlier results and in particular in [19, 21, 24, 4, 8], decidability is ensured only when

urgency is completely disallowed, or places are all bounded. Further, our framework is powerful enough to capture systems of timed finite state machines communicating through *bag channels* [12, 13], with urgency, throughput and latency characteristics of channels, and still yields decidability results.

Structure of the paper. Section 2 introduces Timed-Arc Petri Nets with Urgency (Timed-Arc PNU), their semantics, and gives examples of communication channels that can be represented with this new model. Section 3 examines decidability issues for Timed-Arc PNU and introduces restrictions for decidability of control-state reachability and boundedness. Section 4 addresses the reachability problem for Timed-Arc Petri Nets and TPNs with restricted urgency and gives the main decidability results. Section 5 provides a proof of the main theorem followed by discussion and the conclusion.

2 Timed-Arc Petri nets with Urgency

We will denote by $\mathbb{Q}_{\geq 0}$ the set of positive rational numbers, and by $\mathcal{I}(\mathbb{Q}_{\geq 0})$ the set of intervals over $\mathbb{Q}_{\geq 0} \cup \{\infty\}$. These intervals can be of the form (a, b) , $(a, b]$, $[a, b)$, or $[a, b]$. We will denote by $\mathbb{M}_{\mathbb{R}}$ the set of *multisets* of positive real numbers. For two multisets A and B , we denote by $A \sqcup B$ the disjoint union of A and B , i.e., the multiset that gathers elements of multisets A and B without deleting identical elements. Similarly, we define $A \setminus B$ as the operation that removes from A exactly one occurrence of each element of B (if it exists).

We introduce our main model, Timed-Arc Petri Nets with urgency constraints. The model is based on a semantics using *timed markings* $m : P \rightarrow \mathbb{M}_{\mathbb{R}}$ which associate to each place a multiset describing the ages of all the tokens in this place.

Definition 1 *A Timed-Arc Petri Net with Urgency, denoted Timed-Arc PNU, is a tuple $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ where*

- P is a set of places, T is a set of transitions, m_0 is the initial timed marking,
- $\bullet() : T \rightarrow P$ and $()^\bullet : T \rightarrow P$ are respectively, the backward and forward flow relations indicating tokens consumed/produced by each transition.
- $\gamma : P \times T \rightarrow \mathcal{I}(\mathbb{Q}_{\geq 0})$ is a set of token-age constraints on arcs and
- $U : T \rightarrow \mathbb{Q}_{\geq 0} \cup \{\infty\}$ is a set of urgency constraints on transitions.

For a given arc constraint $\gamma(p, t) = [\alpha(p, t), \beta(p, t)]$ we will call $\alpha(p, t)$ the lower bound and $\beta(p, t)$ the upper bound of $\gamma(p, t)$. Such constraints mean that the transition t is enabled when for each place p of its preset $\bullet t$, there is a token in p of age in $\gamma(p, t)$, i.e., between $\alpha(p, t)$ and $\beta(p, t)$. The urgency constraint $U(t)$ means that a transition must fire if t has been enabled (by its preset of tokens) for $U(t)$ units of time. A Timed-Arc Petri Net [2] can be seen as a Timed-Arc PNU with $U(t) = \infty$ for all $t \in T$. Note that we do not label transitions, hence each transition can be seen as labeled by its unique name.

As an example, consider the Timed-Arc PNU \mathcal{N}_1 of Figure 1. Places are represented by circles, transitions by narrow rectangles, and flow relations by arcs between places and transitions. Urgency of a transition is represented below the transition (in the example, transition t_3 has urgency 3). Arc constraints γ are represented as intervals below

arcs. When unspecified, an arc constraint is set to $[0, \infty)$ and an urgency constraint to ∞ (e.g. $U(t_2) = \infty$). Intuitively, Figure 1 depicts a process $p1$ that sends an unbounded number of messages to a process $p2$ through a channel. A message is sent at least every five time units (t.u.) because of the urgency constraint on t_1 . Latency (or delay) for each message is at least 2 t.u. before being received, and the maximal throughput (or rate) of the channel is between 1 message every t.u. and 1 message every 4 t.u. Changing constraint $[2, \infty)$ into $[2, 100]$ models message loss, i.e., messages not received after 100 t.u. are considered lost. *Formal Semantics of Timed-Arc PNU*: We now define the

semantics of a Timed-Arc PNU $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ in terms of timed markings and discrete and timed moves. For a given place p and timed marking m , we will let age_p denote real values from $m(p)$ depicting the age of one token in place p . Note that as $m(p)$ is a multiset, two tokens in a place p may have identical ages.

We say that a transition t is *enabled* from a timed marking m if, for each $p \in \bullet t$, there exists $age_p \in m(p)$ such that $age_p \in \gamma(p, t)$. A transition t is said to be *urgent* from a timed marking m if $\forall p \in \bullet t, \exists age_p \in m(p)$ such that $\alpha(p, t) + U(t) \leq age_p \leq \beta(p, t)$, i.e., if the preset of t has tokens at least $U(t)$ time units older than required by $\gamma(p, t)$. Let t be an urgent transition from m . This implies that t is enabled. Further, as formally defined below, presence of urgent transitions disallows time from elapsing. Thus, there will exist a place $p \in \bullet t$ such that the oldest token $age_p \in m(p)$ with $age_p \leq \beta(p, t)$ will satisfy $age_p = \alpha(p, t) + U(t)$. An urgent transition t will force occurrence of a discrete move, but not necessarily of this transition t as several transitions can be enabled (or even urgent) at the same time. Formally, the semantics of Timed-Arc PNU is decomposed into timed moves and discrete moves.

Timed moves symbolize elapsing of δ time units from a timed marking in the following way: for a given timed marking m , we denote by $m + \delta$ the timed marking obtained by adding δ to the age of every token: if $m(p) = \{age_1, \dots, age_k\}$, then $(m + \delta)(p) = \{age_1 + \delta, \dots, age_k + \delta\}$. A *timed move* of $\delta > 0$ time units is allowed from m if for every $0 \leq \delta' < \delta$, the timed marking $m + \delta'$ has no urgent transition, and we denote $m \xrightarrow{\delta} m + \delta$ such timed moves.




Fig. 1. Timed-Arc Petri Net with Urgency \mathcal{N}_1 .

Discrete moves represent firings of transitions from a marking m . One can fire transition t from marking m and reach marking m' , denoted $m \xrightarrow{t} m'$ iff t is enabled and for each place p , we have $m'(p) = (m(p) \setminus S_p) \sqcup S'_p$, where

- $S_p = \{age_p\}$ where $age_p \in m(p) \cap \gamma(p, t)$ if $p \in \bullet t$, and $S_p = \emptyset$ otherwise.
- $S'_p = \{0\}$ if $p \in t^\bullet$, and $S'_p = \emptyset$ otherwise.

A Timed-Arc PNU \mathcal{N} defines a timed transition system $\llbracket \mathcal{N} \rrbracket$ whose states are timed markings and transitions are discrete and timed moves. A *run* of \mathcal{N} is a sequence $m_1 a_1 m_2 \cdots m_n$ where, for all $i \in \{1, \dots, n\}$, m_i is a timed marking of \mathcal{N} and $a_i \in (\mathbb{R}_{>0} \cup T)$, such that $m_i \xrightarrow{a_i} m_{i+1}$ is a timed ($a_i \in \mathbb{R}_{>0}$) or discrete ($a_i \in T$) move. We will denote by $\text{Reach}(\mathcal{N})$ the set of reachable timed markings of \mathcal{N} (starting from m_0). An (*untimed*) marking is a function from P to \mathbb{N} . For a timed marking m , we will denote by $m^\# : P \rightarrow \mathbb{N}$ the untimed marking that associates to every place $p \in P$ the number of tokens in $m(p)$. A place $p \in P$ of a Timed-Arc PNU is called *bounded* if there exists an integer K such that for every timed marking $m \in \text{Reach}(\mathcal{N})$, $m^\#(p) \leq K$ and \mathcal{N} is *bounded* if all its places are bounded.

3 Undecidable and Decidable Problems for Timed-Arc PNU

In this paper we will tackle the decidability of the following problems:

- *Reachability*: given a Timed-Arc PNU \mathcal{N} , given an (untimed) marking m , does there exist a timed marking $m' \in \text{Reach}(\mathcal{N})$ with $m'^\# = m$?
- *Control State reachability* (also called *place-reachability*): given a Timed-Arc PNU \mathcal{N} and a place p , does there exist $m \in \text{Reach}(\mathcal{N})$ with $m^\#(p) \geq 1$?
- *Boundedness*: given a Timed-Arc PNU \mathcal{N} , does there exist K such that for all $m \in \text{Reach}(\mathcal{N})$, we have $m^\#(p) \leq K$ for all places p ?

Proposition 1. *Control State reachability, Reachability and Boundedness are undecidable for Timed-Arc PNU.*

Proof (sketch). Reachability is undecidable for Timed-Arc PNU since it is already undecidable for Timed-Arc Petri nets [25]. Because of urgency, control state reachability and boundedness are also undecidable for Timed-Arc PNU. As the proofs closely follow the proofs of undecidability for TPNs [20], we do not detail them here (see also [19] for the proof for Timed-Arc Petri Nets with age invariants). \square

To obtain decidability, two main approaches have been explored. The first involves dropping all urgency requirements. For Timed-Arc PNU, doing so we get back Timed-Arc Petri Nets and their decidability results. For TPNs, this corresponds to the weak semantics [24], under which the reachable (untimed) markings are the markings reachable by the associated (untimed) Petri net. The second approach considers only bounded nets [15] (see also bounded TPNs [7]). Our goal in this paper is to define restrictions for Timed-Arc PNU that ensure decidability for models combining urgency and unbounded nets. This allows us to verify networks of timed systems communicating via unbounded channels with specified latency and throughput such as the one shown in Fig. 1.

3.1 Restricted Urgency

We now define our subclass of Timed-Arc PNU. Let $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ be a Timed-Arc PNU. We start by defining the untimed Petri net associated with \mathcal{N} as $\overline{\mathcal{N}} = (P, T, \bullet(), ()^\bullet, m_0)$, by just dropping the timing constraints. We also define the restriction of \mathcal{N} to a subset of places $P_b \subseteq P$ as the Timed-Arc PNU $\mathcal{N}_{P_b} = (P_b, T, \bullet()_{P_b}, ()^\bullet_{P_b}, m'_{0P_b}, \gamma_{P_b}, U)$, where $\bullet()_{P_b}, ()^\bullet_{P_b}, m'_{0P_b}, \gamma_{P_b}$ are respectively restriction of $\bullet(), ()^\bullet, m_0, \gamma$ to $P_b \times T, T \times P_b$ and P_b . For a timed marking m of \mathcal{N} , we define the timed marking m_{P_b} of \mathcal{N}_{P_b} with $m_{P_b}(p) = m(p)$ for all $p \in P_b$. We observe that if the places that are projected away do not use urgency then every run of \mathcal{N} is also a run in the projected net \mathcal{N}_{P_b} . Formally,

Lemma 1 *Assume that for each transition $t \in T$ with $U(t) < \infty$, we have $\bullet t \subseteq P_b$. Then for every run $m_0 a_1 \dots a_n m_n$ of \mathcal{N} , $m_{0,P_b} a_1 \dots a_n m_{n,P_b}$ is a run of \mathcal{N}_{P_b} .*

Proof. Let $\rho = m_0 a_1 \dots a_n m_n$ be a run of \mathcal{N} . Consider, for the sake of contradiction, the smallest i for which, we do not have $m_{i-1,P_b} \xrightarrow{a_i} m_{i,P_b}$. The only possibility is to contradict urgency because discrete moves satisfy $m_{i-1,P_b} \xrightarrow{a_i} m_{i,P_b}$. Hence we must have $a_i = \delta$ and for some $\delta' < \delta$, there exists a transition t with urgency $U(t) = k < \infty$ and for all input places $p \in P_b$ of t , there is at least one token of age at least k in $(m_{P_b,i} + \delta')(p)$. Now, by assumption, $\bullet t \subseteq P_b$. That is, there are no other places (from P_u) that can be in the preset of t . Hence, by definition of urgency, $m_i \xrightarrow{\delta} m_{i+1}$ is not a valid timed move in \mathcal{N} either, which is a contradiction. \square

Note that the converse is not true in general: a run of \mathcal{N}_{P_b} needs not be a run of \mathcal{N} . We can now define our decidable subclass of Timed-Arc PNU. It is mainly based on the notion of *restricted urgency*, which intuitively means that urgency can be enforced only on the bounded part of the system.

Definition 2 *A Timed-Arc Petri Net with restricted Urgency (denoted Timed-Arc PNrU) is a triple (\mathcal{N}, P_u, P_b) , where $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ is a Timed-Arc PNU, and $P_u \sqcup P_b = P$ is a partition of places of \mathcal{N} such that:*

- for each transition $t \in T$ with $U(t) < \infty$, we have $\bullet t \subseteq P_b$ and,
- the (untimed) Petri Net $\overline{\mathcal{N}_{P_b}}$ associated with \mathcal{N}_{P_b} is bounded.

Intuitively, in a Timed-Arc PNrU (\mathcal{N}, P_u, P_b) , urgency can only be used by transitions consuming tokens from structurally bounded places. As an example, consider the net \mathcal{N}_1 from Fig 1. Let $P_b = \{p1, \text{wait}, \text{throughput}\}$ and $P_u = \{\text{latency}, p2\}$ be a partition of places in P . The (unbounded) places in P_u do not use urgency and the (untimed) Petri Net $\overline{\mathcal{N}_{P_b}}$ is a 1-bounded Petri Net. Hence \mathcal{N}_1 is a Timed-Arc PNrU.

Checking membership in Timed-Arc PNrU, i.e., checking whether a Timed-Arc PNU is with restricted Urgency, is decidable. This immediately follows from the fact that it is decidable whether a place of a Petri Net is bounded. Given \mathcal{N} , it suffices to define $P_b = \bigcup_{U(t) < \infty} \bullet t$, and to check that $\overline{\mathcal{N}_{P_b}}$ is a bounded (untimed) Petri Net. Though we will often refer to places in P_u as “unbounded places”, this only means the contents of these places can be unbounded, not that they must be. On the other hand, places of P_b are bounded in \mathcal{N} :

Lemma 2 *Let (\mathcal{N}, P_u, P_b) be a Timed-Arc PNrU. Then every $p \in P_b$ is bounded in \mathcal{N} .*

Proof. Let K be the bound on the number of tokens in $\overline{\mathcal{N}_{P_b}}$. For the sake of contradiction, if $p \in P_b$ was not bounded in \mathcal{N} , there would exist a reachable marking m with more than K tokens in p . Let $m_0 a_1 \cdots a_n m_n$ be a run reaching $m_n = m$. Then by Lemma 1, $m_{0, P_b} a_1 \cdots a_n m_{n, P_b}$ is a run of \mathcal{N}_{P_b} , and thus of $\overline{\mathcal{N}_{P_b}}$, and m_{n, P_b} has more than K tokens in p , a contradiction with $\overline{\mathcal{N}_{P_b}}$ being K bounded. \square

We next turn to the (un)decidable properties for this subclass of Timed-Arc PNrU.

Theorem 1. *Control-State reachability and Boundedness are decidable for Timed-Arc PNrU. However, reachability is undecidable for Timed-Arc PNrU.*

Proof (sketch). The decidability of control-state reachability and boundedness for Timed-Arc PNU is adapted from [2, 1], by defining a well-quasi order over the markings and using the theory of well structured transition systems [17]. The well quasi order \preceq is defined in the following way. First, we define a region abstraction for markings of Timed-Arc PNrU. This abstraction is a combination of regions of a finite timed automaton representing the behavior of the net on its bounded part, and regions representing symbolically the markings of the unbounded places of the net. This set of regions is equipped with a comparison relation \preceq that requires equality on the region bounded part, and comparable contents on the unbounded part. This relation is compatible with markings comparison and is a well-quasi order. We can then define a successor relation among regions that is an abstract representation of moves of a Timed-Arc PNrU. Regions equipped with their ordering and this successor relation form a well-structured transition system and hence control-state reachability and boundedness are decidable. Details are omitted as the construction is rather similar to [2, 1]. The undecidability of reachability for Timed-arc PNrUs follows directly from the undecidability of reachability for Timed-Arc Petri nets [25, 1]. \square

We remark that the results of the above Theorem 1 can easily be extended to a strictly larger class of Timed-Arc PNUs, where \mathcal{N}_{P_b} is bounded instead of $\overline{\mathcal{N}_{P_b}}$. However, checking membership in this extended class is not decidable as boundedness is not decidable for Timed-Arc PNU.

4 Decidability of the Reachability Problem

In this section we tackle the decidability of the reachability problem. On one hand, reachability is undecidable for Timed-Arc Petri Nets [25], and thus for Timed-Arc PN(r)Us, because an unbounded number of clocks can be encoded, one for each token. On the other hand, Timed-transition Petri Nets (TPNs) [22] only use a bounded number of clocks (one per transition), even if the places have unboundedly many tokens. Nevertheless, (unrestricted) urgency makes reachability undecidable for TPNs [20]. To obtain decidability of reachability, we thus consider classes of TPNs with restricted urgency.

4.1 Timed-Transition Petri Nets (TPNs)

Timed-transition Petri Nets (TPNs for short), also called Time Petri Nets, introduced in [22], associate time intervals to transitions of a Petri net. Formally, a TPN \mathcal{N} is a tuple $(P, T, \bullet(\cdot), (\cdot)^\bullet, m_0, I)$ where P is a finite set of *places*, T is a finite set of *transitions*, $\bullet(\cdot), (\cdot)^\bullet : P \rightarrow T$ are the *backward* and *forward* flow relations respectively, $m_0 \in \mathbb{N}^P$ is the *initial* (untimed) marking, and $I : T \mapsto \mathcal{I}(\mathbb{Q}_{\geq 0})$ maps each transition to a *firing interval*. We denote by $A(t)$ (resp. $B(t)$) the lower bound (resp. the upper bound) of interval $I(t)$. A *configuration* of a TPN is a pair (m, ν) , where m is an untimed *marking* (recall that in untimed markings, $m(p)$ is the number of tokens in p), and $\nu : T \rightarrow \mathbb{R}_{\geq 0}$ associates a real value to each transition. 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 valuation ν associates to each enabled transition $t \in En(m)$ the amount of time that has elapsed since this transition was last newly enabled. An enabled transition t is *urgent* if $\nu(t) \geq B(t)$, with $B(t)$ the upper bound of $I(t)$. An example of a TPN is depicted in Figure 2 below.

We first recall the intermediate marking semantics [7, 5] for TPNs defined using timed and discrete moves between configurations. A *timed move* consists of letting time elapse in a configuration. For (m, ν) , $\nu + \delta$ is defined by $\nu + \delta(t) = \nu(t) + \delta$, for all $t \in En(m)$. A timed move from (m, ν) to $(m, \nu + \delta)$, denoted $(m, \nu) \xrightarrow{\delta} (m, \nu + \delta)$, is allowed if for every $0 \leq \delta' < \delta$, the configuration $(m, \nu + \delta')$ has no urgent transition. A *discrete move* consists of firing an enabled transition t that has been enabled for a duration that fulfills the time constraint attached to t . We have $(m, \nu) \xrightarrow{t} (m', \nu')$ if $t \in En(m)$, $\nu(t) \in I(t)$ and $m' = m - \bullet t + t^\bullet$, for ν' defined below. We call *intermediate marking* the marking $m - \bullet t$ which is obtained after t consumes tokens from its preset but did not create new ones yet. We will say that a transition $t' \in En(m')$ is *newly enabled* by firing of t if either $t' = t$, or $t' \notin En(m - \bullet t)$, i.e. is not enabled in the intermediate marking $m - \bullet t$. Now, we define $\nu'(tt) = 0$ if tt is newly enabled, and $\nu'(tt) = \nu(tt)$ for all $tt \in En(m)$ but not newly enabled. That is, for a transition t both consuming and producing a token in p having a single token, a transition t' with $p \in \bullet t'$ is disabled then newly enabled when t is fired.

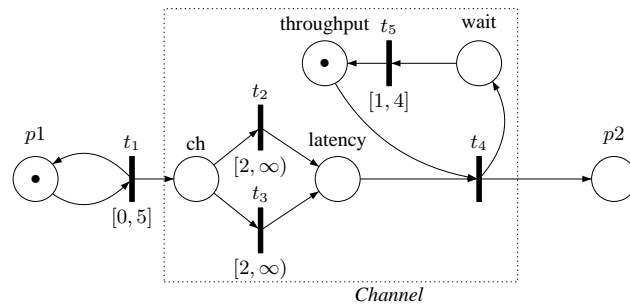


Fig. 2. A TPN \mathcal{N}_2 .

This classical semantics of TPN is somewhat similar to that of Timed-Arc PNU, but is based on configurations instead of timed markings. The only continuous values kept in the configuration of a TPN are in ν . Hence, only $|T|$ “clock” values are kept, and configurations cannot keep track of the exact time elapsed since their creation for arbitrary number of tokens. In particular, a TPN under the intermediate marking semantics cannot encode latency for an unbounded number of tokens [9] (for instance, the property that tokens are consumed at least 2 units of time after each of them is created). More generally, it is not simple to model a channel with specified latency and throughput with the intermediate marking semantics of TPNs. For instance, TPN \mathcal{N}_2 in Figure 2 seems to model a channel with a latency of 2 time units and throughput (rate) of at most 1 message per time unit. However, if a token reaches place ch at date 0 and another at date 1.9, then both can be consumed at time 2, though only one of them has spent two time units at ch , hence it does not faithfully encode a latency of 2.

4.2 A new Timed-Arc Semantics for TPNs

We now introduce a new *timed-arc semantics* in order to model channels with latency and throughput, presenting TPNs as Timed-arc PNUs. The core idea is that the timed-arc semantics takes into account the age of tokens in input places. Formally, we define $\text{Timed}(\mathcal{N})$, the Timed-Arc PNU associated with the TPN $\mathcal{N} = (P, T, \bullet(\cdot), (\cdot)^\bullet, m_0, I)$. Intuitively, $\text{Timed}(\mathcal{N})$ preserves all places and transitions of \mathcal{N} , adds one place p_t per transition t , adds p_t to the pre and post flow of t , and adapts the timing constraints. Figures 3,4 display TPNs $\mathcal{N}_3, \mathcal{N}_4$ on the left and $\text{Timed}(\mathcal{N}_3), \text{Timed}(\mathcal{N}_4)$ on the right. We define $\text{Timed}(\mathcal{N}) = (P', T, \star(\cdot), (\cdot)^\star, m'_0, \gamma, U)$ where:

- $P' = P \cup P_T$ with $P_T = \{p_t \mid t \in T\}$.
- $\star(\cdot), (\cdot)^\star$ extend respectively $\bullet(\cdot), (\cdot)^\bullet$ in the following way: $p \in \star t$ iff $p = p_t$ or $p \in \bullet t$ and $p \in t^\star$ iff $p = p_t$ or $p \in t^\bullet$.
- For all t , for $I(t) = [A(t), B(t)]$, we let $U(t) = B(t) - A(t)$ and for all $p \in \star t$, we set $\gamma(p, t) = [A(t), +\infty)$ (for $I(t) = (A(t), B(t))$ we let $\gamma(p, t) = (A(t), +\infty)$),
- We let $m'_0(p) = 0^{m_0(p)}$ for all $p \in P$ and $m'_0(p_t) = \{0\}$ for all transitions t .

TPN \mathcal{N}_2 under the timed-arc semantics, i.e., $\text{Timed}(\mathcal{N}_2)$, represents the channel with latency 2 and maximal throughput of 1 message per time unit, which is also modeled by the Timed-arc PNU of Figure 1. Indeed, a token can be sent from place ch to place $latency$ by either transition only when it is at least 2 time units old, preserving the latency requirement.

The new timed-arc semantics is close in spirit to the time-on-token semantics of [11], which was defined for 1-safe TPNs. In case of 1-safe TPNs as well as in examples




Fig. 3. A TPN \mathcal{N}_3 (left) which is timed bisimilar to $\text{Timed}(\mathcal{N}_3)$ (right).




Fig. 4. A TPN \mathcal{N}_4 (left) which is not bisimilar to $\text{Timed}(\mathcal{N}_4)$ (right).

such as net \mathcal{N}_3 , this semantics is bisimilar to the classical intermediate marking semantics [11]. However, in general, the behaviors of \mathcal{N} and $\text{Timed}(\mathcal{N})$ differ. Consider for instance the TPN \mathcal{N}_4 in Figure 4. Consider the execution of \mathcal{N}_4 where t_0 fires twice: first at date 0 and then at date 1. At date 2, both t_1 and t_2 have been enabled for 2 time units ($\nu(t_1) = \nu(t_2) = 2$), hence any one of them can fire. Let t_1 fire. Now, t_1 cannot fire again immediately as it is newly enabled (hence $\nu'(t_1) = 0$), but t_2 can fire immediately after t_1 , because $\nu'(t_2) = 2$ (in particular, it is not newly enabled by firing t_1 as there are two tokens in the input place p , i.e., $m(p) - \bullet t(p) = 2 - 1 = 1$).

In contrast, in $\text{Timed}(\mathcal{N}_4)$, if t_0 is fired at date 0 and again at date 1, then at date 2, $m(p) = \{1, 2\}$, and any one of t_1 or t_2 can fire. Just as in the execution of \mathcal{N}_4 , let t_1 fire. After this firing of t_1 , the other transition t_2 cannot fire because $m'(p) = \{1\}$ and $1 < 2$. It is only at date 3 that t_2 can fire. At date 3, transition t_1 cannot fire because $m''(p_{t_1}) = \{1\}$, and $1 < 2$. This illustrates that the behaviors of \mathcal{N}_4 and $\text{Timed}(\mathcal{N}_4)$ can indeed differ in general. In the following, we will use the timed-marking semantics in the general case, where it does not coincide with the intermediate marking semantics.

4.3 Reachability for TPNs under the timed-arc semantics.

Reachability is undecidable for general TPNs because of unrestricted urgency [20], under the timed-arc or the intermediate marking semantics. We now introduce two natural restrictions to urgency to allow decidability.

Definition 3 Let $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, I)$ be a TPN and $P = P_u \sqcup P_b$ be a partition of its places such that the (untimed) Petri Net $\overline{\mathcal{N}}_{P_b}$ associated with \mathcal{N}_{P_b} is bounded.

- \mathcal{N} is called a TPN with restricted urgency if for each transition $t \in T$ with an upper bound $B(t) < \infty$ on its firing interval, we have $\bullet t \subseteq P_b$.
- \mathcal{N} is called a TPN with restricted constraints if for each transition $t \in T$ with a non trivial firing interval $I(t) \neq [0, \infty)$ we have $\bullet t \subseteq P_b$.

The class of TPNs with restricted constraints is strictly contained in the class of TPNs with restricted urgency. As an example, the TPN \mathcal{N}_2 on Fig. 2 is a TPN with restricted urgency but not a TPN with restricted constraints, since there is an arc from the unbounded place ch to transition t_2 with constraints $[2, \infty)$, i.e., a constraint with non-trivial lower bound and no upper (urgency) bound. As for Timed-Arc PNU, checking

whether a TPN is with restricted urgency or with restricted constraints is decidable, since checking boundedness of (untimed) Petri Nets is decidable.

Now, if \mathcal{N} is a TPN with restricted Urgency, then $\text{Timed}(\mathcal{N})$ is a Timed-Arc PNrU, ensuring that boundedness and control-state reachability are decidable. We can now state our main result, namely Theorem 2: reachability is decidable for TPNs with restricted urgency *under timed-arc semantics* (e.g. $\text{Timed}(\mathcal{N}_2)$ from Figure 2 is in this class). TPNs with restricted urgency *under timed-arc semantics* can model networks of (finite-state) *timed* systems with unrestricted urgency, communicating through bag channels [12, 13], specifying maximal throughput and minimal latency, assuming that the throughput is not infinite. Indeed, it suffices to modify the TPN in Figure 2 with $\lceil \frac{x}{\delta} \rceil$ transitions from *ch* to *latency* in order to model a channel with latency at least x and throughput at most δ messages per unit of time.

Theorem 2. *Let \mathcal{N} be a TPN with restricted urgency. Then the reachability, boundedness and control-state reachability problems are decidable for $\text{Timed}(\mathcal{N})$.*

The next section is devoted to the proof of Theorem 2. In essence, we show that although the *timed-arc semantics* of TPNs “formally” uses an unbounded number of clocks, a complex reduction allows to consider only a bounded number of clocks. This step is crucial in the proof of Theorem 2, and we believe that this technique can be generalized and re-used for other problems in related contexts.

5 Proof of Theorem 2

Let (\mathcal{N}, P, Q) , with $\mathcal{N} = (P \cup Q, T, \bullet(\cdot), (\cdot)^\bullet, m_0, I)$ be a TPN with restricted Urgency, P (resp. Q) the set of bounded (resp. unbounded) places. In this section, we show how to check if a given (untimed) marking is reachable in $\text{Timed}(\mathcal{N})$. The intuitive idea is that, under restricted urgency, a transition t which has an unbounded place from Q in its preset, has no urgency/upper constraint. Hence to fire t , it suffices to check the lower bound constraint, i.e., to check that some tokens (among an unbounded number) in its pre-places are old enough. Now, the crucial point is that to check this lower-bound, we need the ages of only a bounded number of tokens, as there are a finite number of transitions, and for each transition t , its associated “clock” p_t is reset after it is fired.

Formally, the proof (of Theorem 2) is in two steps: we first convert the TPN with restricted urgency \mathcal{N} to a TPN with restricted constraints \mathcal{N}' such that $\text{Timed}(\mathcal{N})$ and $\text{Timed}(\mathcal{N}')$ have the same set of reachable markings. In the second step, we obtain a Petri Net that is bisimilar to $\text{Timed}(\mathcal{N}')$, which implies the decidability of reachability.

Step 1: Construction of the TPN with restricted constraints \mathcal{N}' . In order to obtain a TPN with restricted constraints \mathcal{N}' from \mathcal{N} , we will keep (an overapproximation of) ages for a bounded number of tokens from each unbounded place $p \in Q$. For that, we will use $|T| \times |Q|$ gadgets $(C_p^t)_{t \in T, p \in Q}$.

Gadget C_p^t , associated with place $p \in Q$ and transition $t \in T$ (with $p \in \bullet t$), is a TPN with restricted constraints. Each gadget is similar: it has 2 places, 0_p^t and 1_p^t , and in the initial marking the token is at 0_p^t . There is an associated transition $start_p^t$: we have $^*start_p^t = \{p, 0_p^t\}$ and $start_p^{t*} = \{1_p^t\}$, with the timing constraint




Fig. 5. Step 1 of proof: converting (part of) TPN with restricted urgency \mathcal{N} to restricted constraints \mathcal{N}'

$I'(start_p^t) = [0, \infty)$. That is, \mathcal{N}' will non-deterministically guess the transition that will fire. The gadget for a fixed transition t and place p is shown in Figure 5. Every transition t reading from an (unbounded) place $p \in Q$ is transformed to read from (bounded) place 1_p^t of gadget C_p^t instead. That is, if a transition t reads from unbounded places $\{p_1, \dots, p_k\} = \bullet t \cap Q$, then we have $\bullet t = \bullet t \setminus \{p_1, \dots, p_k\} \cup \{1_{p_j}^t \mid j \leq k\}$ and $t^* = t^* \cup \{0_{p_j}^t \mid j \leq k\}$. The timing constraint is left unchanged: $I'(t) = I(t)$. We obtain $\mathcal{N}' = (P', T', \bullet, \bullet^*, m'_0, I')$:

- $P' = P \cup Q \cup \{0_p^t, 1_p^t \mid p \in P, t \in T\}$,
- $T' = T \cup \{start_p^t \mid t \in T, p \in Q\}$,
- \bullet, \bullet^*, I' as defined above, and
- $m'_0(p) = m_0(p)$ for $p \in P$, and $m'_0(0_p^t) = 1, m'_0(1_p^t) = 0$ for all t, p .

It is clear that \mathcal{N}' is a TPN with restricted constraints, with the same set $Q' = Q$ of unbounded places as for all t' with $\bullet t' \cap Q' \neq \emptyset$, we have $t' = start_p^t$ for some $t \in T$, and thus $I(t') = [0, \infty)$.

The idea of the gadget is the following. Let $m \in \text{Reach}(\text{Timed}(\mathcal{N}))$, t be a transition with $I(t) = [a, \infty)$ and $p \in Q \cap \bullet t$ be an unbounded place. $m(p_t)$ is the time elapsed since the last firing of t . For firing t , we need to have both $m(p_t) \geq a$ and $age_p \geq a$, i.e., we need $\min(m(p_t), age_p) \geq a$. In other words, keeping $\min(m(p_t), age_p)$ instead of age_p is sufficient to know whether t is enabled. This is implemented in $\text{Timed}(\mathcal{N}')$, as there can be only one token in 1_p^t , and its age $m'(1_p^t)$ is never older than $m'(p_t)$, as $start_p^t$ can happen only after t fired (0_p^t filled when t fired).

We now show that $\text{Timed}(\mathcal{N}')$ preserves the set of reachable untimed markings of $\text{Timed}(\mathcal{N})$. We start by defining a map f from untimed markings of $\text{Timed}(\mathcal{N}')$ to untimed markings of $\text{Timed}(\mathcal{N})$. Recall that for a timed marking m , m^\sharp refers to the untimed marking obtained by counting the number of tokens in each place. Let $m' \in \text{Reach}(\text{Timed}(\mathcal{N}'))$. For each place $p \in P \cup Q$ of $\text{Timed}(\mathcal{N})$, we define:

$$f(m'^{\sharp})(p) = \begin{cases} m'^{\sharp}(p) + \sum_{t \in T} m'^{\sharp}(1_p^t) & \text{if } p \in Q \\ m'^{\sharp}(p) & \text{otherwise} \end{cases}$$

First, we show that $\text{Timed}(\mathcal{N}')$ can reach only untimed markings corresponding to untimed markings of $\text{Timed}(\mathcal{N})$:

Lemma 3 *Let m' be a timed marking in $\text{Reach}(\text{Timed}(\mathcal{N}'))$. Then there exists a timed marking $m \in \text{Reach}(\text{Timed}(\mathcal{N}))$ with $f(m'^{\sharp}) = m^{\sharp}$.*

Proof. We will prove by induction on the length of a path that if one can reach m' in $\text{Timed}(\mathcal{N}')$, then one can reach m in $\text{Timed}(\mathcal{N})$ such that:

1. for all $p \in P \cup \{p_t \mid t \in T\}$, $m(p) = m'(p)$,
2. for all $q \in Q$, letting T_q be the set of $t \in T$ such that $m'(1_q^t) \neq \emptyset$, we have $m(q) = m'(q) \sqcup \{age_1, \dots, age_k\}$ and there exists a bijection $g : T_q \mapsto [1, k]$ with $m'(1_q^t) \leq age_{g(t)}$ for all $t \in T_q$.

When these two conditions are met, we say that m satisfies the hypothesis wrt m' . It is easy to see that $f(m'^{\sharp}) = m^{\sharp}$ whenever m satisfies the hypothesis wrt m' .

For $m' = m'_0$, we have trivially that m_0 satisfies the hypothesis wrt m'_0 . We can now proceed by induction on the length of run needed to reach m' . Let m' be a reachable marking of $\text{Timed}(\mathcal{N}')$. A path reaching m' ends with a move $m'' \xrightarrow{e} m'$, where e can be a timed move, a firing of a transition $start_p^t$, or a firing of a transition t of the original net. Hence, m'' is reached in less steps than m' . We can hence apply the induction hypothesis, i.e., one can reach m in $\text{Timed}(\mathcal{N})$ with:

1. for all $p \in P \cup \{p_t \mid t \in T\}$, $m(p) = m''(p)$,
2. for all $q \in Q$, letting T_q be the set of $t \in T$ such that $m''(1_q^t) \neq \emptyset$, we have $m(q) = m''(q) \sqcup \{age_1, \dots, age_k\}$ and there exists a bijection $g : T_q \mapsto [1, k]$ with $m''(1_q^t) \leq age_{g(t)}$ for all $t \in T_q$.

For a given bijection g , we denote by $g''(1_q^t) = age_{g(t)}$ the function that relates tokens in places of the form 1_q^t with token $age_{g(t)}$ in m .

Case $e = start_q^t$: We know that m satisfies the hypothesis wrt m'' by hypothesis and want to show that m satisfies the hypothesis wrt m' as well. The conditions are true for all $p \notin \{q, 0_q^t, 1_q^t\}$, as for these places, $m'(p) = m''(p)$. Last, $m(q) = m''(q) \sqcup \{age_1, \dots, age_k\}$ and we have a bijection $g : T_q \mapsto [1, k]$. Now, we have $m'(q) = m''(q) \sqcup \{age_0\}$ for the token age_0 which is consumed by $start_q^t$ from q . Hence $m(q) = m'(q) \sqcup \{age_0, age_1, \dots, age_k\}$, and one can extend $g : T_q \mapsto [1, k]$ to $g' : T_q \cup \{t_q\} \mapsto [0, k]$ by setting $g'(t) = 0$. As $m'(1_q^t) = 0$, we indeed have $age_0 \geq m'(1_q^t)$. Hence m satisfies the hypothesis wrt m' .

Case $e = \delta$ (time elapses by δ units): We note that urgency is not violated in m'' by elapsing δ units of time. Now, since for all $p \in P$ we have $m(p) = m''(p)$, and transitions leaving (unbounded) places of Q have no urgency, urgency is not violated in m either by elapsing δ units of time. Thus $m + \delta$ is reachable in $\text{Timed}(\mathcal{N})$. Finally, it is easy to see that $m + \delta$ satisfies the hypothesis wrt $m' = m'' + \delta$.

Case $e = t$ for some $t \in T$: If $\bullet t$ has only bounded places, since $m(p) = m''(p)$ for all bounded places p , one can fire t from m to obtain a marking m^+ which satisfies the hypothesis wrt m' and we are done. Else, $\bullet t \cap Q \neq \emptyset$ and $I(t) = [a, +\infty)$ for some $a \in \mathbb{R}_{\geq 0}$. For all $q \in \bullet t \cap Q$, we have $m''(1_q^t) \geq a$ as t can be fired from m'' . Taking the token $g''(1_q^t)$ of $m(q)$, we have $g''(1_q^t) \geq m''(1_q^t) \geq a$ for each $q \in \bullet t \cap Q$. Further, since $m(p) = m''(p)$ for all $p \in P \cup \{p_t \mid t \in T\}$, we have that t is enabled from m . We now carefully define a particular marking m^+ of $\text{Timed}(\mathcal{N})$ which can be obtained from m by firing t . First, for every $q \in \bullet t \cap Q$, we delete the token $g''(1_q^t)$ of $m(q)$. Then for all $p \in \bullet t \cap (P \cup \{p_t \mid t \in T\})$, we define age_p the age of token removed from $m''(p)$ to $m'(p)$, and remove it from $m(p)$ as well. Finally, for every place p of t^\bullet , we create a token of age 0 in $m(p)$. It is now easy to check that m^+ satisfies the hypothesis wrt m' . \square

Next, we show that every untimed marking of $\text{Timed}(\mathcal{N})$ can be simulated in $\text{Timed}(\mathcal{N}')$:

Lemma 4 *Let m be a timed marking in $\text{Reach}(\text{Timed}(\mathcal{N}))$. Then one can reach in $\text{Timed}(\mathcal{N}')$ any timed marking m' with:*

- (1) for all $p \in P \cup \{p_t \mid t \in T\}$, we have $m'(p) = m(p)$ and
- (2) for all $t \in T$, $p \in Q$, we have either $m'(0_p^t) = \emptyset$ or $m'(0_p^t) = m'(p_t)$, and
- (3) for all $q \in Q$, letting $T'_q = \{t \in T \mid m'(1_q^t) \neq \emptyset\}$, we have $m(q) = m'(q) \sqcup \{age_t \mid t \in T'_q\}$ with $m'(1_q^t) = \min(m(p_t), age_t)$ for all $t \in T'_q$.

Proof (sketch). We proceed by induction on the length of run reaching m in \mathcal{N} . For a run of length 0, this is trivial. Assume that m is reached after a move $m^- \xrightarrow{e} m$. Let m' be any marking satisfying the conditions (1–3) above wrt m . We will show that we can reach m' in $\text{Timed}(\mathcal{N}')$.

Assume that e is a timed move that lets $\delta > 0$ units of time elapse. Hence, for every place $p \in P \cup Q$, and every token $age_p \in m(p)$, $age_p \geq \delta$. We have $m^- = m - \delta$. We first show that for all $p' \in P'$ and all $age'_p \in m'(p')$, $age'_p \geq \delta$. This is easy to see for $p' \in P \cup \{p_t \mid t \in T\}$ as $m(p') = m'(p')$, and hence also for $p' \in \{0_p^t \mid t \in T, p \in Q\}$. For $p \in Q$, we have $m'(q) \sqsubseteq m(q)$. Further, for all $t \in T$, $p \in Q$ with $m'(1_p^t) \neq \emptyset$, we have $m'(1_p^t) = \{\min(m(p_t), age_p)\}$ with $age_p \in m(p)$. As $m(p_t) \geq \delta$ and $age_p \geq \delta$, we have $m'(1_p^t) \geq \delta$. Thus, $age'_p \geq \delta$ for all $p' \in P'$ and all $age'_p \in m'(p')$. We can now define the timed marking $m'' = m' - \delta$. It is then easy to check that m'' satisfies the conditions (1–3) above wrt to m^- , and so, we can apply the induction hypothesis and conclude that m'' is reachable in $\text{Timed}(\mathcal{N}')$.

Now, we show that waiting δ units of time from m'' is allowed in $\text{Timed}(\mathcal{N}')$. That is, we show that it does not violate any urgency: Suppose not, i.e., suppose the urgency of some transition t was violated. Then, this would imply that $I(t) = [a, b]$ and thus $\bullet t \subseteq P$ contains only bounded places. As m^- and m'' coincide on bounded places, this would also violate urgency on m^- , a contradiction with a δ timed move being allowed from m^- . Hence δ units of time can elapse from m'' , reaching marking m' . Thus m' is reachable in $\text{Timed}(\mathcal{N}')$.

Finally, the case of a discrete move e firing a transition t can be handled by a similar analysis, which completes the proof of this lemma. \square

Now observe that for all places p , we have $\sum_{t \in T} m'^{\sharp}(1_p^t) \leq |T|$. Thus fixing an untimed marking c , there exist only a finite number of untimed markings m'^{\sharp} such that $f(m'^{\sharp}) = c$. Combining Lemmas 3, 4, we obtain:

Proposition 2. *Let c be an untimed marking of $\text{Timed}(\mathcal{N})$. Let c' be any untimed marking of $\text{Timed}(\mathcal{N}')$ with $f(c') = c$. Then c is reachable in $\text{Timed}(\mathcal{N})$ iff c' is reachable in $\text{Timed}(\mathcal{N}')$.*

This completes the first step of the proof of Theorem 2.

Step 2: From the TPN with restricted constraints \mathcal{N}' to a Petri Net \mathcal{N}'' . Now we show that for a TPN with restricted constraints \mathcal{N}' , it is decidable whether a marking c' is reachable in $\text{Timed}(\mathcal{N}')$, by reducing \mathcal{N}' to an equivalent (untimed) Petri net. As marking reachability is decidable for Petri nets, this completes the proof of Theorem 2.

Proposition 3. *For any TPN with restricted constraints \mathcal{N}' , one can construct a Petri Net \mathcal{N}'' such that \mathcal{N}'' and $\text{Timed}(\mathcal{N}')$ are (untimed) bisimilar.*

Proof. Given a TPN with restricted constraints \mathcal{N}' , we first construct a 1-bounded (untimed) Petri Net \mathcal{N}_1 which is bisimilar to $\text{Timed}(\mathcal{N}'_B)$, where \mathcal{N}'_B is the bounded part of \mathcal{N}' . Formally, we have the following lemma.

Lemma 5 *If \mathcal{N}'_B is a K -bounded TPN, for some positive integer K , we can construct a 1-bounded Petri Net \mathcal{N}_1 such that \mathcal{N}_1 and $\text{Timed}(\mathcal{N}'_B)$ are (untimed) bisimilar.*

Proof (sketch). The proof of this lemma is easily obtained by building a timed automaton bisimilar to $\text{Timed}(\mathcal{N}'_B)$ and interpreting its regions as places of a 1-safe Petri Net [19], adapting a result for the intermediate semantics of TPNs [14].

After building the Petri net $\mathcal{N}_1 = (P_1, T_1, \bullet(\cdot), (\cdot)^\bullet, m_1^0)$ we add the unbounded places of \mathcal{N}' . Formally, we build the Petri net $\mathcal{N}_2 = (P_2, T_1, \star(\cdot), (\cdot)^\star, m_2^0)$ with:

- The set P_2 of places of \mathcal{N}_2 is $P_2 = P_1 \cup P_u$, for P_u the unbounded places of $\text{Timed}(\mathcal{N}')$.
- Initial marking m_2^0 is the union of m_1^0 and of the restriction of the initial marking of $\text{Timed}(\mathcal{N}')$ to its set P_u of unbounded places.
- The set of transitions of \mathcal{N}_2 is the set T_1 of transitions of \mathcal{N}_1 . Concerning the flow relations, for $t_1 \in T_1$ and its corresponding transition $t \in T$ in the original net $\text{Timed}(\mathcal{N}')$, we have $p \in \star t_1$ if:
 - $p \in P_1$ and $p \in \bullet t_1$ (arc from p to t_1 in \mathcal{N}_1), or
 - $p \in P_u$ and there is an arc from p to t in $\text{Timed}(\mathcal{N}')$.

We have $p \in t_1^\star$ if $p \in P_1$ and $p \in t_1^\bullet$, or if $p \in P_u$ and there is an arc from t to p in $\text{Timed}(\mathcal{N}')$.

With this, we have the following lemma:

Lemma 6 *$\text{Timed}(\mathcal{N}')$ and \mathcal{N}_2 are (untimed) bisimilar.*

Proof. A timed marking m of $\text{Timed}(\mathcal{N}')$ can be decomposed as $m = m_b \cup m_u$, where m_b is the restriction of m to bounded places, and m_u the restriction to unbounded places. Similarly, a marking of \mathcal{N}_2 can be decomposed as $m_2 = m_1 \cup m'_u$ by restriction to bounded and unbounded places respectively. From Lemma 5 and from the construction of \mathcal{N}' , we know that $\text{Timed}(\mathcal{N}_B)$ is bisimilar to \mathcal{N}' . Let $R_{B,1}$ be the unique largest bisimulation between timed markings of $\text{Timed}(\mathcal{N}_B)$ and markings of \mathcal{N}' .

We denote by R a relation from timed markings of $\text{Timed}(\mathcal{N})$ to markings of \mathcal{N}_2 defined as follows. Let $m = m_b \cup m_u$ be a marking of $\text{Timed}(\mathcal{N})$ and $m_2 = m_1 \cup m'_u$ be a marking of \mathcal{N}_2 . Then, $(m, m_2) \in R$ iff $(m_b, m_1) \in R_{B,1}$, and $m'_u = m_u^\sharp$. Obviously, we have $(m_0, m_2^0) \in R$. We can now prove that R is a bisimulation.

Let $(m, m_2) \in R$. Assume that $m \xrightarrow{\delta} m + \delta \xrightarrow{t} m'$ in \mathcal{N} . Thus $m_b \xrightarrow{\delta} m_b + \delta \xrightarrow{t} m'_b$ in \mathcal{N}_b with m'_b the bounded part of m' . Furthermore, $m_u^\sharp \geq \bullet t \cap P_u$. Thus we have $m_1 \xrightarrow{t} m'_1$ in \mathcal{N}_1 , and furthermore, $(m'_1, m'_b) \in R_{B,1}$. By definition of \mathcal{N}_2 , firing t results in a flow of tokens among places of P_u that is identical (regardless of ages) in \mathcal{N} and in \mathcal{N}_2 , so we indeed have $m_1 \cup m_u^\sharp \xrightarrow{t} m'_1 \cup m'_u$. Furthermore $m'_u = m'_u$, so $(m', m'_2) \in R$.

Conversely, assume that $m_2 \xrightarrow{t} m'_2$. We denote $m_2 = m_1 \cup m_3$ and $m'_2 = m'_1 \cup m'_3$ where m_3, m'_3 denote respectively the projections of m_2 and m'_2 on P_u . In particular, as t can fire, we have $m_1 \xrightarrow{t} m'_1$. So, there exists a reachable marking m'_b of $\text{Timed}(\mathcal{N}_B)$ such that $(m'_b, m'_1) \in R_{B,1}$ and there exists δ such that $m_b \xrightarrow{\delta} m_b + \delta \xrightarrow{t} m'_b$. In particular, δ does not violate any urgency constraints in the bounded part of the net.

Now, \mathcal{N}' is a TPN with restricted constraints. This means that all urgency constraints are in the bounded part of \mathcal{N} . Hence, $m \xrightarrow{\delta} m + \delta$ does not violate any urgency constraints. Now, to show that R is a bisimulation, we want to show that $m + \delta \xrightarrow{t} m'$ is possible in \mathcal{N}' , for some m'_u with $m' = m'_b \cup m'_u$ and $(m'_u)^\sharp = m'_3$. To see this, we start by noting that, since $(m, m_2) \in R$, with $m_2 = m_1 \cup m_3$, we have $m^\sharp(p) = m_3(p) \geq 1$ for all $p \in P_u \cap \bullet t$. Also, we have trivially that $m^\sharp(p) = m_b^\sharp(p) \geq 1$ for all $p \in P_B \cap \bullet t$ as t is enabled from m_1 , and $(m_b, m_1) \in R_{B,1}$. Thus t is enabled. Now, $m + \delta$ respects all the timings constraints of t : as \mathcal{N}' is a TPN with restricted constraints, all constraints apply to the bounded part. Transition t is enabled from m_1 , thus t can fire from $m + \delta$. For the unbounded part, firing of t can consume any token in places of $P_u \cap \bullet t$ and we easily get $(m'_u)^\sharp = m'_3$. For the bounded part, we choose to consume the tokens consumed during the transition $m_b + \delta \xrightarrow{t} m'_b$. We thus obtain $m' = m'_b + m'_u$, and $(m', m'_2) \in R$. Hence R is a bisimulation relation. \square

We can then conclude that the net \mathcal{N}_2 , as constructed above, is bisimilar to $\text{Timed}(\mathcal{N}')$ and hence satisfies the properties required by the proposition. Thus, setting \mathcal{N}_2 to be the net \mathcal{N}'' , we obtain the proof of Proposition 3. \square

From Proposition 2, we have that for every TPN \mathcal{N} with restricted urgency, one can build (Step 1) a TPN \mathcal{N}' with restricted constraints that has the same set of reachable markings. Then proposition 3 shows that one can design (Step 2) a Petri net that is bisimilar to \mathcal{N}' . As reachability is decidable for Petri nets, this allows to conclude the proof of Theorem 2.

5.1 Discussion

Let us now observe some salient points regarding the proof of Theorem 2 and in particular, how it relies on several features of the considered nets. First, the proof works only for nets with restricted urgency. If urgency is not restricted, one can easily model unbounded counters with places, and zeros tests with urgency, which yields undecidability of reachability, control-state reachability and boundedness. Second, Step 1 of the proof of Theorem 2 works only with a timed-arc semantics. The main idea in this step was to simulate clocks with gadgets as in Figure 5, that need to be assembled to obtain nets with restricted constraints, which are equivalent (i.e., have the same set of behaviors). However, for TPNs under the intermediate semantics, assembling the gadgets leads to nets that are not equivalent.

Step 2 of the proof works for both the intermediate and the timed-arc semantics. Thus, starting from a TPN with restricted constraints, we get decidability of reachability for TPNs with intermediate semantics as stated in the following Theorem 3. However, as seen earlier, this class does not allow to model channels with latency constraints.

Theorem 3. *Let \mathcal{N} be a TPN with restricted constraints. Then the reachability, boundedness and control-state reachability problems are decidable for \mathcal{N} .*

Proof (sketch). The proof of Theorem 3 is obtained by a simple adaptation of Proposition 3 from Section 5, which shows that for any TPN with restricted constraints \mathcal{N} , one can construct a Petri Net \mathcal{N}' that is (untimed) bisimilar. \square

Finally, our proof works only when the considered systems can be implemented with a bounded number of clocks (in order to get a bounded number of gadgets in the proof). This approach would not work for systems modeling channels with latency *and* unbounded throughput, which require nets with an unbounded number of clocks to be specified. Decidability of reachability for such classes is left open.

6 Conclusion

In this paper, we considered extensions of Timed-Arc Petri Nets and subclasses of TPNs to express urgency and latency constraints, while obtaining decidability results for unbounded systems. Decidability is obtained when urgency is used only in the bounded part of the system. This led us to consider a timed-arc semantics for general TPNs, defined via a Timed-Arc Petri Nets with Urgency. This new timed-arc semantics allows TPNs to model restricted forms of latency, namely, unbounded latency in a channel can be modeled when the throughput of this channel is bounded, as well as urgency requirements. Further, the new timed-arc semantics is also interesting as decidability of reachability can be proved for a class of TPNs larger with the timed-arc semantics than with the intermediate marking semantics. Table 1 in the Introduction summarizes the decidability results as well as expressiveness in terms of which (subclasses of) models allow latency and urgency. The *relative expressiveness* of classes (wrt timed bisimilarity) is summarized in Figure 6, where we also emphasize their decidability status. While Timed-Arc Petri nets are contained in Timed-Arc Petri nets with restricted Urgency, Timed Arc Petri nets (with restricted Urgency) and TPNs are disjoint classes of




Fig. 6. Inclusion of classes of Time/Timed-Arc Petri Nets with urgency w.r.t. timed bisimilarity.

models. For the entire subclass of Timed-Arc Petri nets with restricted Urgency, reachability is decidable. However, outside this class, i.e. without restriction on the use of urgency, control-state reachability and boundedness are undecidable. Further, with a timed-arc semantics, TPNs fall back into the class of Timed-Arcs PNs with Urgency. And by restricting urgency and under this timed-arc semantics, we obtain decidability of reachability. Further, as a subclass of Timed-Arcs PNs with restricted Urgency, control-state reachability and boundedness are decidable. Finally, the class of TPNs under their intermediate marking semantics does not enjoy decidability results. However, by restricting to the class of TPNs with restricted constraints, one gets decidability of reachability and control-state reachability. The decidability of reachability for TPNs with restricted urgency under intermediate marking semantics remains open.

As future work, we plan to study robustness properties, i.e, whether the system can withstand infinitesimal timing errors, as has been extensively studied for timed automata [23, 16, 10], etc. We would like to extend the study started for TPNs (e.g. [3]) to Timed-Arc Petri Nets with restricted Urgency.

Acknowledgments: We thank the anonymous reviewers for their helpful suggestions which led to significant improvements in the presentation of this article. This work was partially supported by Indo-French CEFIPRA project AVERTS, by DST-INSPIRE faculty award [IFA12-MA-17], and by the DISTOL associated team of INRIA.

References

1. P. A. Abdulla, P. Mahata, and R. Mayr. Dense-timed Petri nets: Checking Zenoness, token liveness and boundedness. *Journal of Logical Methods in Computer Science*, 3(1), 2007.
2. P. A. Abdulla and A. Nylén. Timed Petri nets and BQOs. In *ICATPN, LNCS 2075*, p.53-70, 2001.

3. S. Akshay, L. Hélouët, C. Jard, and P.-A. Reynier. Robustness of time Petri nets under guard enlargement. In *RP, LNCS 7550*, pages 92–106, 2012.
4. F. Bause and P. S. Kritzinger. *Stochastic Petri nets - an introduction to the theory (2. ed.)*. Vieweg, 2002.
5. 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, 2005.
6. B. Bérard, F. Cassez, S. Haddad, D. Lime, and O. H. Roux. The expressive power of time petri nets. *TCS*, 474:1–20, 2013.
7. B. Berthomieu and M. Diaz. Modeling and verification of time dependent systems using time Petri nets. *IEEE Trans. in Software Engineering*, 17(3):259–273, 1991.
8. B. Berthomieu, F. Peres, and F. Vernadat. Bridging the gap between timed automata and bounded time petri nets. In *FORMATS'06*, volume 4202 of *LNCS*, pages 82–97, 2006.
9. H. Boucheneb, D. Lime, and O.H. Roux. On multi-enabledness in time Petri nets. In *ICATPN*, volume 7927 of *LNCS*, pages 130–149, 2013.
10. P. Bouyer, N. Markey, and O. Sankur. Robustness in timed automata. In *Proc. of RP*, volume 8169 of *LNCS*, pages 1–18, 2013.
11. Th. Chatain and C. Jard. Back in Time Petri Nets. In *FORMATS*, volume 8053 of *LNCS*, pages 91–105, 2013.
12. L. Clemente, F. Herbreteau, A. Stainer, and G. Sutre. Reachability of communicating timed processes. In *FoSSaCS*, volume 7794 of *LNCS*, pages 81–96, 2013.
13. L. Clemente, F. Herbreteau, and G. Sutre. Decidable topologies for communicating automata with fifo and bag channels. In *CONCUR, LNCS 8704*, p. 281-296, 2014.
14. D. D'Aprile, S. Donatelli, A. Sangnier, and J. Sproston. From time Petri nets to timed automata: An untimed approach. In *TACAS, LNCS 4424*, pages 216–230, 2007.
15. A. David, L. Jacobsen, M. Jacobsen, and J. Srba. A forward reachability algorithm for bounded timed-arc Petri nets. In *SSV'12*, volume 102 of *EPTCS*, pages 125–140, 2012.
16. M. De Wulf, L. Doyen, N. Markey, and J.-F. Raskin. Robust safety of timed automata. *Formal Methods in System Design*, 33(1-3):45–84, 2008.
17. A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! In *TCS*, 256(1-2):63–92, 2001.
18. S. Haddad. Time and timed Petri nets. In *Disc PhD School 2011*, pages available at <http://www.lsv.ens-cachan.fr/haddad/disc11-part1.pdf>, 2011.
19. L. Jacobsen, M. Jacobsen, M. H. Møller, and J. Srba. Verification of timed-arc Petri nets. In *SOFSEM'11*. LNCS 6543, p.46-72, 2011.
20. N. D. Jones, L. H. Landweber, and Y. E. Lien. Complexity of some problems in Petri nets. *Theor. Comput. Sci.*, 4(3):277–299, 1977.
21. J. A. Mateo, J. Srba, and M. G. Sørensen. Soundness of timed-arc workflow nets. In *PETRI NETS'14*. LNCS 8489, p.51-70, 2014.
22. P. M. Merlin. *A Study of the Recoverability of Computing Systems*. PhD thesis, University of California, Irvine, CA, USA, 1974.
23. A. Puri. Dynamical properties of timed automata. In *DEDS*, 10(1-2):87–113, 2000.
24. P.-A. Reynier and A. Sangnier. Weak time Petri nets strike back! In *Proc. of CONCUR 2009*, volume 5710 of *LNCS*, pages 557–571, 2009.
25. V. V. Ruiz, F. C. Gomez, and D. de Frutos-Escrig. On non-decidability of reachability for timed-arc Petri nets. In *Pnpm*, pages 188–. IEEE Computer Society, 1999.
26. B. Walter. Timed Petri-Nets for Modelling and Analysing Protocols with Real-Time Characteristics. In *Proc. of PSTV*, pages 149–159, 1983.