



Timed Petri Nets with (restricted) Urgency

Sundararaman Akshay, Blaise Genest, Loïc Hélouët

► **To cite this version:**

Sundararaman Akshay, Blaise Genest, Loïc Hélouët. Timed Petri Nets with (restricted) Urgency. 2014. <hal-01088997>

HAL Id: hal-01088997

<https://hal.inria.fr/hal-01088997>

Submitted on 29 Nov 2014

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.

Timed Petri Nets with (restricted) Urgency

S. Akshay¹, Blaise Genest², Loïc Hélouët³

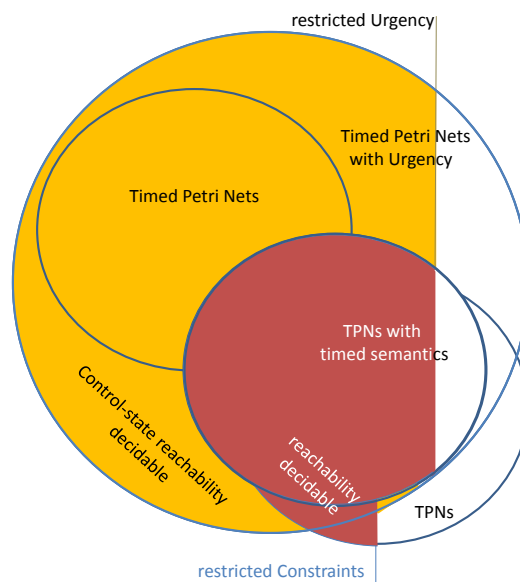
¹ IIT Bombay, India

² CNRS, Team SUMO, UMR IRISA, Rennes, France

³ INRIA, Team SUMO, Rennes, France

Abstract. Time Petri Nets (TPN) [Mer74] and Timed Petri Nets [Wal83] are two incomparable classes of concurrent models with timing constraints: urgency cannot be expressed using Timed Petri Nets, while TPNs can only keep track of a bounded number of continuous values (“clocks”). We introduce Timed Petri Nets with Urgency, extending Timed Petri Nets with the main features of TPNs.

We present upto-our-knowledge the first decidability results for Petri Net variants combining time, urgency and unbounded places. First, we obtain decidability of *control-state reachability* for the subclass of Timed Petri Nets with Urgency where urgency constraints can only be used on bounded places. By restricting this class to use a finite number of “clocks”, we further show decidability of (marking) *reachability*. Formally, this class corresponds to TPNs under a new, yet natural, *timed semantics* where urgency constraints are restricted to bounded places. Further, under their original semantics, we obtain the decidability of *reachability* for a more restricted class of TPNs.



1 Introduction

Petri nets are a simple yet powerful mathematical formalism for modeling distributed systems. In order to specify real-time behaviors, they can be enriched with quantitative information in the form of timing constraints. There are several such extensions of Petri Nets. We first discuss the features and short-comings of two main variants, namely: *Time Petri Nets (TPNs)* [Mer74] and *Timed Petri Nets* [Wal83,AN01].

TPNs can constrain each transition with a timing interval. To be fireable, a transition needs to have been enabled for an amount of time inside the given interval. Also, 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, termination, control-state reachability, boundedness) are undecidable for TPNs [JLL77], as two counter machines can easily be encoded. To obtain decidability, one either restricts to bounded TPN [BD91], or gives up urgency [RS09]. In this 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.

On the other hand, Timed Petri Nets associate (continuous, positive valued) ages to each token. 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 Petri Nets cannot encode urgency [Had11,AN01]. Although the number of token ages is unbounded, control-state reachability and boundedness are decidable for Timed Petri Nets [AN01]. The reason is that without urgency, we obtain monotonicity properties for tokens (as they can be “lost” by just staying forever at a place), which permits the application of the theory of well structured transition systems [FS01]. However, reachability is still undecidable for Timed Petri Nets [RGE99].

In terms of expressivity, Timed Petri Nets cannot express that a token is produced exactly every unit of time: due to lack of urgency, it is not possible to force a transition to fire. On the other hand, TPNs cannot express that an unbounded number of tokens (with slightly different ages) are consumed at least two units of time after they have been created. We now propose a formalism which can easily specify these two characteristics (see Fig. 1), and more. Namely, we introduce *Timed Petri nets with Urgency*, extending Timed Petri Nets [Wal83,AN01] with explicit urgency requirements, à la Merlin [Mer74]. This is done by introducing, in Timed Petri Nets, urgency constraints on transitions, forcing the transition to fire if it remains enabled for long enough.

Unsurprisingly, most problems are undecidable because of urgency (Proposition 2), as is the case for TPNs. To get around this, we consider a class of systems where urgency can be used, but is restricted in a meaningful way. In this class of nets, transitions consuming tokens exclusively from bounded places can use urgency; other transitions that consume tokens from at least one unbounded place do not have urgency constraint. We say that such systems (be it Timed Petri Nets or TPN) are *with restricted Urgency*. Notice that it is in general *not* the case that the untimed language of a TPN with restricted Urgency is the language of the associated Petri Net (without the timing constraints).

From a practical point of view, this allows to specify e.g. a timed system of finite state machines communicating through bag channels [CHS14]: each process is a finite-state machine – hence can use urgency. However, (unbounded) channels can use (latency) constraints on tokens but not urgency: a message from a channel cannot be forced to be received in a given time bound (but it can be lost after a time out period). This extends the channel models of [CHSS13], where urgency is not allowed.

We present upto-our-knowledge the first decidability results for such Petri Net variants combining time, urgency and unbounded places. For Timed Petri Nets with restricted Urgency, we obtain decidability of *control-state reachability* (Theorem 1), and more generally of timed marking coverability, boundedness, termination, etc. The proof is based on the fact that we can transform such a system into an equivalent Timed Petri Net. The construction separates the bounded and unbounded part, processing the bounded part to remove urgency thanks to the finite number of states, and then reinserts the unbounded part (which does not use urgency, hence is already a Timed Petri Net).

Undecidability of the *reachability* problem for Timed Petri Nets with restricted Urgency, inherited from Timed Petri Nets, is due to the presence of unboundedly many timed tokens. We thus want to define a subclass of Timed Petri Nets with Urgency where the number of (useful) timed tokens is bounded. This leads us to consider TPNs, which inherently use a bounded number of “clocks”. We define an alternative *timed semantics* for TPNs, presenting them as a subclass of Timed Petri Nets with Urgency. We then obtain our main result: a class of Timed Petri Nets with restricted Urgency such that *reachability* is decidable (Theorem 4). This class corresponds to TPNs with restricted urgency under the timed semantics.

The timed semantics is close to the original semantics, in the sense that both semantics coincide for a large class of TPNs (Proposition 1). Under their *standard* semantics, we obtain decidability of reachability only for a strict subclass of TPNs with restricted urgency (Theorem 3), showing the algorithmic advantage of the timed semantics.

The paper is organized as follows: Section 2 introduces Timed Petri Nets with Urgency and their semantics. Section 3 compares this model with TPNs and define their timed semantics. Section 4 introduces the reachability and boundedness problems that we consider and shows that these are all undecidable in general for Timed Petri Nets with Urgency. Section 5 defines restricted urgency subclasses for Timed Petri nets with Urgency and for TPNs and states the associated decidability results. Section 6 gives a sketch of the proofs, whose details can be found in the appendix.

2 Timed Petri nets with Urgency

We will denote by $\mathbb{Q}_{>0}$ the set of positive rational numbers, and by $\mathcal{I}(\mathbb{Q}_{>0})$ the set of intervals over $\mathbb{Q}_{>0} \cup \infty$. 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).




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

We introduce our main model, Timed Petri Net with urgency constraints. The model is based on a *timed 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 Petri net with Urgency*, denoted *Timed 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,
- $\bullet() : T \rightarrow P$ is a backward flow relation indicating tokens consumed by each transition,
- $()^\bullet : T \rightarrow P$ is a forward flow relation indicating tokens produced by each transition,
- m_0 is the initial timed marking,
- $\gamma : P \times T \rightarrow \mathcal{I}(\mathbb{Q}_{\geq 0})$ is a set of token-age constraints on arcs,
- $U : T \rightarrow \mathbb{Q}_{\geq 0}$ 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 it has been enabled (by its preset of tokens) for $U(t)$ units of time. A Timed Petri Net [AN01] can be seen as a Timed PNU with $U(t) = +\infty$ for all $t \in T$. Notice that we do not label transitions, as we are interested in properties on markings rather than on languages. Each transition can be seen as labeled by its unique name. When building equivalent models for a given net \mathcal{N} , we will have to duplicate transitions. In this case, every transition of the new net will be associated to a transition (name) of \mathcal{N} in an obvious way.

As an example, consider the Timed PNU \mathcal{N}_1 of Figure 1. Places are represented by circles, transitions by narrow rectangles, and flow relation by arcs between places and transitions. Urgency of a transition t is represented below the transition (in the example, transition t_1 has urgency 0). Arc constraints γ are represented as intervals below arcs. When unspecified, an arc constraint is set to $[0, +\infty)$ and an urgency constraint to $+\infty$ (e.g. $U(t_2) = +\infty$). Intuitively, this net represents process r sending messages to process s via channel c . Process r sends messages with bandwidth of one message per unit of time. The messages reach s with latency of at least 2. Messages not received after the time out period of 100 units of time are considered “lost”.

Formal Semantics of Timed PNU We now define the semantics of a Timed PNU $\mathcal{N} = (P, T, \bullet(\cdot), (\cdot)^\bullet, m_0, \gamma, U)$ in terms of timed markings and discrete and timed moves. For a given place p and timed marking m , we will denote by $age \in m(p)$ 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 place p may have identical ages.

We will say that a transition t is *enabled* from a timed marking m if and only 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 *urgent* from a timed marking m if $\forall p \in \bullet t, \exists age_p \in m(p)$ such that $\alpha(p, t) + U(t) \leq age_p \leq \beta(p, t)$. Notice that any urgent transition t is enabled. 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. As formally defined below, presence of urgent transitions disallows time from elapsing. Because of this, for every *reachable* urgent transition t , 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 $age_p = \alpha(p, t) + U(t)$. Formally, the semantics of Timed 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 time moves.

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.

In other words, when t fires, a token with age age_p that satisfies the arc-constraint $\gamma(p, t)$ is consumed in each input place of t , and a new token with age 0 is produced in each output place. In particular, tokens that are consumed by firing a transition t are the tokens which enabled the transitions. They need not be unique: several ages from the same place can enable the transition, any of them can be consumed. Hence discrete moves are not a priori deterministic. Further, note that as in TPNs, timed moves are not allowed when a transition is urgent. This ensures that a discrete move must happen when a transition t is urgent. Note however that another transition t' may be fired (t' needs not be urgent). After t' fires, it is possible that no transitions (including t) are urgent anymore (because the corresponding tokens have been consumed), and then time can elapse. Else, urgency remains and a discrete move still needs to happen.

With the above semantics, a Timed PNU \mathcal{N} defines a timed transition system $\llbracket \mathcal{N} \rrbracket$ whose states are timed markings and transitions are discrete and timed moves. 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 $P \rightarrow \mathbb{N}$. For a timed marking m , we will denote by $m^\# : P \rightarrow \mathbb{N}$ the marking that associates to every place $p \in P$ the number of tokens in $m(p)$. We will say that a place $p \in P$ of a Timed PNU is *bounded* if there exists an integer K such that for every timed marking $m \in \text{Reach}(\mathcal{N})$, $m^\#(p) \leq K$. We will say that \mathcal{N} is *bounded* iff there exists a K such that all the places are bounded by K .

Continuing with our channel example of Fig. 1, the initial timed marking m_0 has $m_0(r) = \{0\}$, $m_0(c) = m_0(s) = \emptyset$, i.e., one counter with age 0 in place r . The only moves allowed from m_0 are a sequence of timed moves with maximal duration of 1 time unit, leading to a timed marking m_1 with $m_1(r) = \{1\}$, $m_1(c) = m_1(s) = \emptyset$. In m_1 , transition t_1 is enabled and urgent, hence time cannot elapse further. As t_1 is the only enabled transition, it fires, and marking m_2 is reached with $m_2(r) = \{0\}$, $m_2(c) = \{0\}$, $m_2(s) = \emptyset$. From m_2 , no transition is enabled, but as there is no urgent transition, time can elapse. After letting one unit of time elapse, a timed marking m_3 is reached with $m_3(r) = \{1\}$, $m_3(c) = \{1\}$, $m_3(s) = \emptyset$. From m_3 , t_1 is enabled and urgent, and it fires as the only enabled transition. After t_1 fires, one unit of time must elapse, leading to timed marking m_4 with $m_4(r) = \{1\}$, $m_4(c) = \{1, 2\}$, $m_4(s) = \emptyset$. Both transition t_1 and t_2 are enabled in m_4 . Transition t_1 is urgent (but not transition t_2), hence time cannot elapse. Transition t_2 can fire, and then timed marking m_5 is reached with $m_5(q) = \{1\}$, $m_5(r) = \{1\}$, $m_5(s) = \{0\}$. Then t_1 is still urgent and it is the only transition enabled in m_5 , thus it fires and the run proceeds in this manner.

Example: Imul in Intel Silvermont Architecture (2014 Atom line). Timed PNU can be used to specify models describing systems with time constraints such as latency or throughput. The Intel Optimization Guide specifies that integer multiplication on port 1 of Silvermont Architecture has 3 to 5 units of time of latency (because of pipeline stages) and throughput of one instruction per 1 to 2 unit of time (see also <http://www.realworldtech.com/silvermont/5/>).

The net \mathcal{N}_2 of Fig. 2 models this situation. A token in *imul-inst* models the issue of an integer multiplication instruction. A latency token is created 3 to 5 units of time after a token in *imul-inst*. The throughput token is ready 1 to 2 unit of time after the previous instruction fired (else it is not-ready). The instruction is processed as soon as both latency and throughput places are filled. This consumes one instruction and restarts the throughput process. If too many instructions are scheduled exceeding the throughput, they are stuck in the pipeline till the throughput token is ready, which is what actually happens in a processor. Note that without urgency, an instruction might stay forever in *imul-inst* or in latency.

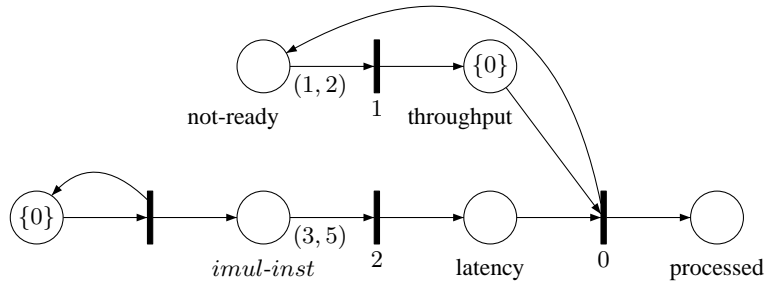


Fig. 2. Timed Petri Net with Urgency \mathcal{N}_2 .

3 Comparison with TPNs

Introduced in [Mer74], *Time Petri nets (TPNs for short)* associate a time interval to each transition of a Petri net. A time Petri net \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*, $\bullet(\cdot) : P \rightarrow T$ is the *backward* flow relation, $(\cdot)^\bullet : P \rightarrow T$ is the *forward* flow relation, $m_0 \in \mathbb{N}^P$ is the *initial* (untimed) marking, and $I : T \mapsto \mathcal{I}(\mathbb{Q}_{\geq 0})$ associates with each transition 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 $m(p)$ is the number of tokens in p), and $\nu : T \rightarrow \mathbb{R}_{\geq 0}$ associates a real-time 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 $I(t)$. We use the classical intermediate marking semantics (see for instance [BD91]) defined as follows, 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 constraints 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 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 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.

We define newly enabledness via intermediate markings since other options (e.g. using the so-called *atomic* or *persistent* semantics [RS09]) lead to undecidability even in the absence of urgency [RS09].




Fig. 3. A (non simple) TPN \mathcal{N}_3 (left) which is not bisimilar to $Timed(\mathcal{N}_3)$ (right).

The semantics of TPN is somewhat similar to that of Timed 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 configuration cannot keep track of the exact time elapsed since their creation for arbitrary number of tokens. Hence no TPN can encode that an unbounded number of tokens are processed at least 2 unit of times after they are created (latency at least 2), which can be easily done using Timed Petri Nets (see our two examples in Fig. 1,2).

For a TPN $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, I)$, we can define $\text{Timed}(\mathcal{N})$ its *associated Timed PNU*. 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. We let $\text{Timed}(\mathcal{N}) = (P', T, {}^*(), ()^\bullet, m'_0, \gamma, U)$ where:

- $P' = P \cup P_T$ with $P_T = \{p_t \mid t \in T\}$.
- ${}^*(), ()^\bullet$ extend respectively $\bullet(), ()^\bullet$ in the following way: $p \in {}^*t$ iff $p = p_t$ or $p \in \bullet t$ and $p \in t^\bullet$ 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 $p \in {}^*t$, 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 .

We display in Fig.3 a TPN \mathcal{N}_3 on the left and $\text{Timed}(\mathcal{N}_3)$ on the right. We show that for a subclass of (possibly unbounded) TPNs, $\text{Timed}(\mathcal{N})$ and \mathcal{N} are timed bisimilar.

Definition 2. A TPN \mathcal{N} is said to be *simple* if for all reachable configurations, for all places p , either p has at most 1 token, or only one transition t is enabled with $p \in \bullet t$.

The TPN \mathcal{N}_3 on Figure 3 is not simple. However, it becomes simple if we remove e.g. transition t_2 . Intuitively, simple TPNs have the property that immediately after firing a transition with input p , the clock of every enabled transition with input p is 0.

Proposition 1. Let \mathcal{N} be a simple TPN. Then $\text{Timed}(\mathcal{N})$ and \mathcal{N} are timed bisimilar.

Proof (sketch). The bisimulation relation \mathcal{R} between reachable configurations of \mathcal{N} and reachable timed markings of $\text{Timed}(\mathcal{N})$ is the following. We associate a timed marking m' of $\text{Timed}(\mathcal{N})$ to each configuration (m, ν) of \mathcal{N} , with $m'^{\sharp}(p) = m(p)$ for all $p \in P$, and such that for all $t \in \text{En}(m)$, $\nu(t) = \min_{p \in {}^*t} (\max'_p)$, with \max'_p the oldest token in $m'(p)$ for all p . Recall that $p_t \in {}^*t$. Notice that we can safely consider \max'_p as by definition $\beta(p, t) = +\infty$ for all p, t .

That is, we consider the oldest tokens *oldest_p* in every place p of the net. Then we consider the youngest *oldest_p* for p in the preset of t . When \mathcal{N} is simple, this token age measures the time elapsed since the transition t became enabled in $\text{Timed}(\mathcal{N})$. This is exactly what is measured by $\nu(t)$ in \mathcal{N} . This is true for all transitions, which implies that (m, ν) of \mathcal{N} and m' of $\text{Timed}(\mathcal{N})$ are bisimilar. A formal proof of this is included in the appendix. \square

Notice however that if \mathcal{N} is not simple, then \mathcal{N} and $\text{Timed}(\mathcal{N})$ may differ. Consider for instance the TPN \mathcal{N}_3 in Figure 3. Consider the execution where the TPN fires t_0 twice: first at date 0 and then at date 1. At date 2, both t_1 and t_2 have been enabled for 2 unit of time ($\nu(t_1) = \nu(t_2) = 2$), hence any one of them can fire. Let say t_1 fires. 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., $m(p) - \bullet t(p) = 2 - 1 = 1$). By contrast, in $\text{Timed}(\mathcal{N})$, 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, let say t_1 . But 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 $1 < 2$.

We now consider $\text{Timed}(\mathcal{N})$ as a new semantics of a TPN \mathcal{N} , called the *timed semantics* of \mathcal{N} . One advantage of this timed semantics with respect to the standard (time) semantics is that firing of transitions consume tokens that have enabled this transition. Further, we will show in the following that the timed semantics enjoy better algorithmic properties: the class of TPN for which reachability is decidable is more restricted for the standard semantics (Theorem 3) than for the timed semantics (Theorem 4).

We do not believe that TPN are included into Timed PNU in general: $\nu(t)$ cannot be computed in general by looking only at age of actual tokens as it may depends on consumed tokens which prevented disabling t . Also, knowing whether a transition is newly enabled corresponds to performing a kind of zero test, hence it seems hard to code it with new transitions and places. As it subsumes the class of Timed Petri Nets, the class of Timed PNUs is not included into the class of TPNs [Had11].

4 Problems Statement and Undecidability

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

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

Another problem that is frequently addressed is the *coverability* question: given a *timed* marking m , is there a marking m' in $\text{Reach}(\mathcal{N})$ such that $m \leq m'$? For Timed Petri Nets (and extensions), it is easy to show that this problem and control state reachability are inter-reducible, hence we do not consider it explicitly.

Because of urgency, every non-trivial problem is undecidable for Timed PNU, following the proof of undecidability for TPNs from [JLL77]:

Proposition 2. *Control State reachability, Reachability and (place) boundedness are undecidable for Timed PNU.*

An unsurprising solution to obtain decidability in TPNs is to drop the urgency requirements [RS09], but then reachable (untimed) markings are just the (untimed) markings reachable by the untimed Petri Net obtained by dropping every timing constraint.

5 Restricted Urgency and Decidability

As general Timed PNU are Turing equivalent, we now introduce subclasses ensuring decidability. We will provide the decidability results in this section, as well as the sequence of lemmas that lead to them. Some proof ideas will be given in the next section, while the details can be found in the appendix.

We start by defining the restriction of a Timed PNU to a subset of places. Let $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ be a timed PNU, and let $Q \subseteq P$ be a subset of places. The restriction of \mathcal{N} to Q is the Timed PNU $\mathcal{N}_Q = (Q, T, \star()_Q, ()^\star_Q, m'_{0Q}, \gamma_Q, U)$, where $\star()_Q, ()^\star_Q, m'_{0Q}, \gamma_Q$ are respectively restriction of $\bullet(), ()^\bullet, m_0, \gamma$ to $Q \times T, T \times Q$ and Q . Note that every run of \mathcal{N} is a run of \mathcal{N}_Q (by projecting markings on Q), but the converse is not true in general. Our main notion is *restricted urgency*, which intuitively means that urgency is allowed only on the bounded part of the system.

Definition 3. A Timed Petri Net with restricted Urgency (*denoted* Timed PNrU) is a Timed PNU $\mathcal{N} = (P, T, \bullet(), ()^\bullet, m_0, \gamma, U)$ together with a partition $P_u \sqcup P_b = P$ of its places such that:

- For each transition $t \in T$ with $\bullet t \cap P_u \neq \emptyset$, we have $U(t) = \infty$.
- The restriction of \mathcal{N} to places of P_b is bounded.

Intuitively, this means that urgency cannot be used for transitions consuming tokens from unbounded places. For instance, Fig. 1 displays a Timed PNrU (only place c is unbounded). This is not the case of the Timed PNU in Fig. 2 (latency and imul-inst are unbounded places red by two transitions with urgency). However, the number of instructions in flight in a processor is actually bounded (around 50), hence it is possible to express the system using a Timed PNrU.

Urgency allows to perform zero test-like operations as shown in the undecidability proof of proposition 2. Forbidding zero testing unbounded places seems reasonable to obtain decidability. Proofs showing decidability for TPNs without urgency (as defined and proved in [RS09]) use the fact that timing constraints can be forgotten everywhere without impacting the untimed language. This is not the case with our restriction as urgency is allowed in some parts of the net.

As each Timed Petri Net is a Timed PNrU (it suffices to choose $P_u = P$), the undecidability of reachability for Timed PNrU is inherited from Timed Petri Nets. We show that we can however decide control state reachability and boundedness.

Theorem 1. *Control-State reachability and (Place) boundedness are decidable for Timed PNrU. However, reachability is not decidable for Timed PNrU.*

The decidability proofs are obtained by encoding any Timed PNrU into an equivalent Timed Petri Net. Hence they have the same set of reachable markings. Now, we immediately obtain decidability from the following result of [AN01]:

Theorem 2. [AN01] *Control-State reachability and (Place) boundedness are decidable for Timed Petri Nets.*

5.1 Decidability of Reachability for a subclass of TPNs

A natural question is whether decidability can be obtained for TPNs as well, by restricting urgency. In fact, a condition stronger than restricted urgency is needed.

Definition 4. Let $\mathcal{N} = (P, T, \bullet(\cdot), (\cdot)^\bullet, m_0, I)$ be a TPN and $P = P_u \sqcup P_b$ be a partition of its places such that the restriction of \mathcal{N} to places of P_b is bounded. Then,

- \mathcal{N} is called a TPN with restricted urgency (TPNrU for short) if for each transition $t \in P_u$ with $\bullet t \cap P_u \neq \emptyset$, we have $I(t) = [c, \infty)$ with $c \in \mathbb{Q}_{\geq 0}$.
- \mathcal{N} is called a TPN with restricted constraints if for each transition $t \in P_u$ with $\bullet t \cap P_u \neq \emptyset$, we have $I(t) = [0, \infty)$.

The class of TPNs with restricted constraints is strictly contained in the class of TP-NrU. As an example, the TPN \mathcal{N}_3 on Fig. 3 has restricted urgency but not restricted constraints, since place p is unbounded and there is a transition with non-trivial lower constraints reading from p , but no transition with urgency constraints.

We show that *reachability* is decidable for TPNs with restricted constraints. This is done by encoding \mathcal{N} into an (untimed) bisimilar Petri Net:

Proposition 3. For each TPN with restricted constraints \mathcal{N} , one can construct a Petri Net \mathcal{N}' such that \mathcal{N}' and \mathcal{N} are bisimilar.

Decidability of reachability for Petri Nets can thus be extended:

Theorem 3. Reachability is decidable for TPNs with restricted constraints.

We do not know if reachability is decidable for the class of TPnrU. However, when we consider TPnrU *under the timed semantics*, we obtain decidability of reachability, as shown now.

5.2 Decidability of Reachability for a subclass of Timed PNrU

Let \mathcal{N} be a TPnrU. Clearly, $\text{Timed}(\mathcal{N})$ is a Timed PNrU. Then:

Theorem 4. Let \mathcal{N} be a TPN with restricted Urgency. Then the reachability problem for $\text{Timed}(\mathcal{N})$ is decidable.

In order to prove Theorem 4, we encode $\text{Timed}(\mathcal{N})$ as a machine $\text{Timed}(\mathcal{N}')$ such that \mathcal{N}' is a TPN with restricted constraints. By adapting Proposition 3 to produce a Petri Net bisimilar to $\text{Timed}(\mathcal{N}')$, we obtain decidability of reachability.

Thus, the set $\{\text{Timed}(\mathcal{N}) \mid \mathcal{N} \text{ is a TPnrU}\}$ is a subclass of Timed PNrU for which we obtain decidability of reachability.

6 Proofs

6.1 Proof of Theorem 1

We start by defining a notion of timed equivalence that we need for this proof. Let S be a system, either a Timed PNU, a Timed Petri Net or a timed automaton [AD94], where each transition is associated to one transition of T . A timed sequence of S is a finite sequence $(t^1, d^1) \dots (t^n, d^n)$ with $t^i \in T$ a transition name and $0 < d^1 < \dots < d^n$ timing instants, such that from the initial state s_0 of the (timed transition) semantics of S , one can perform the sequence of transitions: $s_0 \xrightarrow{d^1} r_1 \xrightarrow{t^1} s_1 \xrightarrow{d^2-d^1} r_2 \dots \xrightarrow{d^n-d^{n-1}} r_n \xrightarrow{t^n} s_n$. Two systems are said to be *timed equivalent* if they have the same set of (finite) timed sequences. A crucial observation is that if two systems are timed bisimilar, they are timed equivalent (the converse does not hold in general).

Now, the proof of Theorem 1 is performed in two steps. First, given a Timed PNrU \mathcal{N} , we construct a 1-bounded (labeled) Timed Petri Net \mathcal{N}_1 which is *timed equivalent* to the bounded part \mathcal{N}_B of \mathcal{N} (it is not possible to get bisimilarity because Timed Petri Nets do not use urgency). This can be done by first converting the bounded part of the Timed PNrU into a bisimilar timed automaton with all states accepting (as done for TPNs e.g., in [DDSS07]), and then using the construction from [BHR08] (Theorem 7 and Corollary 1) to convert the timed automaton to 1-bounded Timed Petri Net which is timed equivalent to it. We denote $\mathcal{N}_1 = (P_1, T_1, \bullet, \cdot, m_1, \gamma_1)$ this 1-bounded Timed Petri Net.

In the second step, we show that the original Timed PNrU \mathcal{N} is timed equivalent to the net \mathcal{N}_2 formed by adding unbounded places of \mathcal{N} to the \mathcal{N}_1 . Formally, we construct the Timed Petri Net $\mathcal{N}_2 = (P_2, T_1, \star, \cdot, m_2, \gamma_2)$ as follows:

- 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 \mathcal{N} .
- Initial timed marking m_2 is the union of m_1 and of the restriction of the initial timed marking of \mathcal{N} restricted to its unbounded places P_u .
- Finally, the set of transitions of \mathcal{N}_2 is the set T_1 of transitions of \mathcal{N}_1 , and the flow relations and γ_2 are defined in the following way. Let $t_1 \in T_1$, corresponding to a transition $t \in T$ in the original net \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), and then $\gamma_2(p, t_1) = \gamma_1(p, t)$, or
 - $p \in P_u$ and there was an arc from p to t in \mathcal{N} , and then $\gamma_2(p, t_1) = \gamma(p, t)$, the constraints from \mathcal{N} .

Similarly, we have $p \in t_1^\star$ if $p \in P_1$ and $p \in t_1^\bullet$ (arc from t_1 to p in \mathcal{N}_1), or if $p \in P_u$ and there is an arc from t to p in \mathcal{N} .

Lemma 1. \mathcal{N} and \mathcal{N}_2 are timed equivalent.

The first step along with Lemma 1 thus completes the proof of Theorem 1.

6.2 Proof of Proposition 3

The proof is similar to the above proof. First, given a TPNrU \mathcal{N} , we construct a 1-bounded (untimed) Petri Net \mathcal{N}_1 which is bisimilar to the bounded part \mathcal{N}_B of \mathcal{N} .

This can be done by using [DDSS07] to build a timed automata bisimilar to \mathcal{N}_B and interpreting its regions as a 1 safe Petri Net. Notice that the exact same process can be done starting with $\text{Timed}(\mathcal{N}_B)$ instead of \mathcal{N}_B . We provide the proof in appendix.

Lemma 2. [DDSS07] *If \mathcal{N} is a K -bounded TPN, we can construct a 1-bounded Petri Net \mathcal{N}_1 such that \mathcal{N}_1 and \mathcal{N} are (untimed) bisimilar.*

Now, the second step shows that the original TPNrU \mathcal{N} is bisimilar to the Petri net obtained by adding unbounded place of \mathcal{N} to the Petri net $\mathcal{N}_1 = (P_1, T_1, \bullet, \cdot, m_1)$. Formally, we construct the Petri net $\mathcal{N}_2 = (P_2, T_1, \bullet, \cdot, m_2)$ as previously:

- 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 \mathcal{N} .
- Initial marking m_2 is the union of m_1 and of the restriction of the initial marking of \mathcal{N} restricted to its unbounded place P_u .
- Finally, the set of transitions of \mathcal{N}_2 is the set T_1 of transitions of \mathcal{N}_1 , and the flow relations are defined in the following way. Let $t_1 \in T_1$ corresponding to a transition $t \in T$ in the original net \mathcal{N} . We have $p \in {}^*t_1$ if:
 - $p \in P_1$ and $p \in \bullet t_1$ (arc from p to e in \mathcal{N}_1), or
 - $p \in P_u$ and there was an arc from p to t in \mathcal{N} .
Similarly, we have $p \in t_1 \cdot$ if $p \in P_1$ and $p \in t_1 \cdot$ (arc from e to p in \mathcal{N}_1), or if $p \in P_u$ and there is an arc from t to p in \mathcal{N} .

Lemma 3. *\mathcal{N} and \mathcal{N}_2 are (untimed) bisimilar.*

The bisimulation relation we use is: (m, ν) is in relation with m_2 iff there exists m_u a marking of P_u and $m = m_b \cup m_u$, $m_2 = m_1 \cup m_u$ with (m_b, ν_b) bisimilar to m_1 according to Lemma 2. Lemma 3 thus completes the proof of proposition 3.

6.3 Proof of Theorem 4

Let (\mathcal{N}, P, Q) , with $\mathcal{N} = (P \cup Q, T, \bullet(\cdot), (\cdot)^\bullet, m_0, I)$ be a TPNrU, where P is the set of bounded places and Q the set of unbounded places. The idea of the proof is the following: intuitively, the timed semantics of \mathcal{N} is that for a transition to fire, each of its token need to be old enough. If the transition has an unbounded place from Q in its preset, then there is no urgency, and this condition is sufficient. Else it is bounded and we can deal with it easily. Assume wlog. that \mathcal{N}_P is 1-bounded. For unbounded places, we need their ages only for a bounded number of them, as there is a finite number of transitions, and each transition time is reset after it is fired.

Construction of the TPN with restricted constraints \mathcal{N}' : In order to obtain a TPN with restricted constraints \mathcal{N}' , we will keep (an overapproximation of) ages for a bounded number of tokens from an unbounded place. Basically, for each unbounded place $p \in Q$, we will create $|T|$ channels $(C_p^t)_{t \in T}$ for this place. Each channel C_p^t from place $p \in Q$ is a TPN with restricted constraints. Each channel is similar: It has 2 places: 0_p^t (that is C_p^t empty) and 1_p^t (one token in C_p^t). The initial marking is 0_p^t . There is an associated transition $start_p^t$: we have ${}^*start_p^t = \{p, 0_p^t\}$ and $start_p^{t \cdot} = \{1_p^t\}$. The timing constraint is $I'(start_p^t) = [0, +\infty)$. That is, \mathcal{N}' will non deterministically guess the transition to start counting for.

Now, we transform every transition to be able to read from any of these channels c_p^i instead of the places p . If a transition t reads from unbounded places $\{p_1, \dots, p_k\} = \bullet t \cap Q$, then we have $\star t = \bullet t \setminus \{p_1, \dots, p_k\} \cup \{1_{p_j}^t \mid j \leq k\}$ and $t^\star = t \bullet \cup \{0_{p_j}^t \mid j \leq k\}$. The timing constraint is left unchanged: $I'(t) = I(t)$.

Now, let $\mathcal{N}' = (P', T', \star(), ()^\star, m'_0, I')$:

- $P' = P \cup Q \cup \{0_p^i, 1_p^i \mid p \in P, i \leq |T|\}$,
- $T' = T \cup \{\text{start}_p^i \mid i \leq |T|, p \in Q\}$,
- $\star(), ()^\star, I'$ as defined above, and
- $m'_0(p) = m_0(p)$ for $p \in P$, and $m'_0(0_p^i) = 1$, $m'_0(1_p^i) = 0$ for all i, p

It is clear that \mathcal{N}' is a TPN with restricted constraints, with $Q' = Q$ the set of unbounded states, as for all t with $\star t \cap Q \neq \emptyset$, we have $t = \text{start}_p^i$, and thus $I(t) = [0, +\infty)$.

Let m' be a marking of $\text{Timed}(\mathcal{N}')$. Define $c = f(m'^\#) : P \cup Q \cup P_T$ such that for all $p \in P \cup Q \cup P_T$, $c(p) = m'^\#(p) + \sum_{i \leq |T|} m'^\#(1_p^i)$ (the second part is ignored for $p \in P \cup P_T$). Notice that because $\sum_{i \leq |T|} m'^\#(1_p^i) \leq |T|$ for all p , fixing c , there exists only a finite number of $m'^\#$ such that $f(m'^\#) = c$. For two timed marking m, m' , we write $m \equiv m'$ whenever $f(m'^\#) = m'^\#$.

Now, remark that because there is only one token possible in 1_p^t and it can be started only after t fired (because of 0_p^t filled when t fired), $m'(1_p^i)$ can never be older than $m'(p_t)$, the age since the last firing of t . Let t be a transition, $I(t) = [a, \infty)$ and let $p \in Q \cap \bullet t$. Now, assume that we have an age $\text{age}_p \in m(p)$ greater than $m(p_t)$. A necessary condition for firing the transition t is both $\tau(p_t) \geq a$ and $\text{age}_p \geq a$. This necessary condition can be summed up as $\min(\tau(p_t), \text{age}_p) \geq a$. It is actually what our construction does: keeping $\min(\tau(p_t), \text{age}_p)$ for some age of a token in p , because it is necessary to know the real age of age_p if older than $\tau(p_t)$. Notice that this condition holds with both the original and the timed semantics of TPNs.

Let c be an untimed marking of $\text{Timed}(\mathcal{N})$. Let c' be any untimed marking of $\text{Timed}(\mathcal{N}')$ with $f(c') = c$. To obtain the proof of Theorem 4, it suffices to apply the following proposition, whose second statement is decidable because \mathcal{N}' is a TPN with restricted constraints.

Proposition 4. *c is reachable in $\text{Timed}(\mathcal{N})$ iff c' is reachable in $\text{Timed}(\mathcal{N}')$.*

The proof uses the two following lemmas relating \mathcal{N} and \mathcal{N}' .

Lemma 4. *Let m' be a reachable marking of $\text{Reach}(\text{Timed}(\mathcal{N}'))$. Then there exists a marking $m \in \text{Reach}(\text{Timed}(\mathcal{N}))$ with $m \equiv m'$.*

Lemma 5. *Let m be a reachable marking of $\text{Reach}(\text{Timed}(\mathcal{N}))$. Then one can reach in $\text{Timed}(\mathcal{N}')$ any m' with:*

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

Lemmas 4 and 5 imply Proposition 4.

7 Conclusion

In this paper, we extend Timed Petri Nets to express urgency constraints, thus capturing the essential feature of TPNs. While general Timed Petri Nets with Urgency are undecidable (as are TPNs), we obtain decidability when urgency is used only in the bounded part of the system. We then consider an alternative timed semantics for general TPNs in terms of Timed Petri Nets with Urgency. Compared to the original semantics, our timed semantics ensures that transitions only consume tokens that enabled them long enough. Also, decidability of reachability can be proved for a larger class of systems.

We plan to study robustness properties, i.e., whether the system can withstand infinitesimal timing errors. Robustness has been extensively studied for timed automata [Pur00,DDMR08,BMS13], etc. We would like to extend the study started for TPNs (e.g. [AHJR12]) to Timed Petri Nets with restricted Urgency.

References

- [AD94] R. Alur and D. Dill. A theory of timed automata. *In TCS*, 126(2):183–235, 1994.
- [AHJR12] 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.
- [AN01] P. Aziz Abdulla and A. Nylén. Timed Petri nets and BQOs. *In ICATPN*, pages 53–70. Springer, 2001.
- [BD91] 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.
- [BHR08] P. Bouyer, S. Haddad, and P.-A. Reynier. Timed Petri nets and timed automata: On the discriminating power of Zeno sequences. *Inf. Comput.*, 206(1):73–107, 2008.
- [BMS13] P. Bouyer, N. Markey, and O. Sankur. Robustness in timed automata. *In Proc. of RP*, volume 8169 of LNCS, pages 1–18. Springer, 2013.
- [CHS14] 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.
- [CHSS13] L. Clemente, F. Herbreteau, A. Stainer, and G. Sutre. Reachability of communicating timed processes. *In FoSSaCS*, volume 7794 of LNCS, pages 81–96. Springer, 2013.
- [DDMR08] Martin De Wulf, Laurent Doyen, Nicolas Markey, and Jean-François Raskin. Robust safety of timed automata. *Formal Methods in System Design*, 33(1-3):45–84, 2008.
- [DDSS07] 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.
- [FS01] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *In TCS*, 256(1-2):63–92, 2001.
- [Had11] 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.
- [JLL77] 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.
- [Mer74] Philip M. Merlin. *A Study of the Recoverability of Computing Systems*. PhD thesis, University of California, Irvine, CA, USA, 1974.
- [Pur00] A. Puri. Dynamical properties of timed automata. *In DEDS*, 10(1-2):87–113, 2000.
- [RGE99] V. V.o Ruiz, F. C. Gomez, and D. de F. Escrig. On non-decidability of reachability for timed-arc Petri nets. *In PNPm*, pages 188–. IEEE Computer Society, 1999.
- [RS09] 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.
- [Wal83] B. Walter. Timed Petri-Nets for Modelling and Analysing Protocols with Real-Time Characteristics. *In Proc. of PSTV*, 1983.

Appendix

We present the detailed proofs below.

Proposition 1. *Let \mathcal{N} be a simple TPN. Then $\text{Timed}(\mathcal{N})$ and \mathcal{N} are timed bisimilar.*

Proof. The bisimulation relation \mathcal{R} between reachable configurations of \mathcal{N} and reachable timed markings of $\text{Timed}(\mathcal{N})$ is the following. We associate a timed marking m' of $\text{Timed}(\mathcal{N})$ to each configuration (m, ν) of \mathcal{N} , with $m'^{\sharp}(p) = m(p)$ for all $p \in P$, and such that (P1) holds:

(P1) for all $t \in \text{En}(m)$, $\nu(t) = \min_{p \in {}^*t}(\max'_p)$, with \max'_p the oldest token in $m'(p)$ for all p . Recall that $p_t \in {}^*t$. Notice that we can safely consider \max'_p as by definition $\beta(p, t) = +\infty$ for all p, t .

We prove that for each pair $((m, \nu), m')$ in \mathcal{R} a timed or discrete move a is allowed from (m, ν) if and only if a is also allowed from m' , and the resulting configurations and timed marking is in \mathcal{R} . Clearly, $((m_0, \nu_0), m'_0) \in \mathcal{R}$.

Take $(m, \nu), m' \in \mathcal{R}$. Thus $m'^{\sharp}(p) = m(p)$ for all $p \in P$ and (P1) holds.

First, assume that the elapse of δ units of time is possible from (m, ν) . It means that no urgent transition is met after $\delta' < \delta$ units of time. In particular, for all $t \in \text{En}(m)$, $\nu(t) + \delta' \leq B(t)$. As $\nu(t) = \min_{p \in {}^*t}(\max'_p)$, there exists $p \in {}^*t$ with $\max'_p + \delta' \leq B(t) = A(t) + (B(t) - A(t)) = \alpha(p, t) + U(t)$, thus t is not urgent in m' either for any δ' . Both moves are allowed, and we have $((m, \nu + \delta), m' + \delta) \in \mathcal{R}$ as $\nu(t) = \min_{p \in {}^*t}(\max'_p)$, with \max'_p the oldest token in $m'(p)$, we have $\nu + \delta(t) = \nu(t) + \delta = \min_{p \in {}^*t}(\max'_p) + \delta = \min_{p \in {}^*t}(\max''_p)$, with \max''_p the oldest token in $m' + \delta(p)$. Symetrically, if a δ time units move is possible from m' , then it is possible from (m, ν) as well and lead to this same bisimilar configuration.

Now we address discrete moves: let t be a transition fireable from configuration (m, ν) , and leading to configuration (m_2, ν_2) . Thus $t \in \text{En}(m)$ and $\nu(t) \in I(t)$. For all $p \in {}^*t$, $\max'_p \geq \nu(t) \geq A(t) = \alpha(t)$, hence t is enabled in m' . When firing t from a timed marking of a Timed PNU, there is no unicity of the set of tokens consumed for each place. Hence, we can obtain several timed markings m'_2 . No matter the chosen m'_2 , we get easily $m'_2{}^{\sharp}(p) = m_2(p)$ for all $p \in P$.

To finish proving that $((m_2, \nu_2), m'_2) \in \mathcal{R}$, we need to show that P1 holds. Let \max''_p denote the maximal value of a token in the timed marking m'_2 . For all $tt \in \text{En}(m_2)$, we have several cases:

- if ${}^*\bullet tt \cap \bullet t \neq \emptyset$, then there exists a place p in ${}^*\bullet t \cap \bullet tt$. If $tt = t$, we have by definition $\nu_2(t) = 0 = \text{age}_t$ and $m_2(p_t) = \{\text{age}_t\}$. If $tt \neq t$ and both tt and t are enabled in m , then by definition of simple TPN, $m(p) = 1$. Thus tt is newly enabled in (m_2, ν_2) , as the intermediate marking reached during firing of t from m satisfies $m(p) - \bullet t(p) = 0$. Thus $\nu_2(tt) = 0$. On the other hand, we also have that $p \in t\bullet$ (else tt would not be in $\text{En}(m_2)$) and $m'_2(p) = \{0\}$. Hence $\max''_p = 0$ and hence $\min_{q \in m'_2}(\max''_q) = 0$, that is P1 holds.
- Else, if $tt \in \text{En}(m)$, we have $\max'_p = \max''_p$ for any place p in *tt , as no token from *tt has been consumed. We hence have $\nu_2(tt) = \nu(tt) = \min_{p \in {}^*tt}(\max'_p) = \min_{p \in {}^*tt}(\max''_p)$, and P1 holds.

- Else, tt is newly enabled, with some $p \in \bullet tt \cap t \bullet$ and $m'_2(p) = \{0\}$. Thus $\nu_2(tt) = 0 = \min_{p \in m'_2}(\max''_p)$ as $\max''_p = m'_2(p)$: P1 holds

Thus we have proved that $((m_2, \nu_2), m'_2) \in \mathcal{R}$.

Last, let transition t be fireable from timed marking m' and leading to timed marking m'_2 . As $m'^{\sharp}(p) = m(p)$ for any place $p \in P$, $t \in \text{En}(m)$. As t is fireable from m' , there exists $\text{age}'_p \in m'(p)$ with $\alpha(p, t) \leq \text{age}'_p \leq \max'_p$. Hence $\alpha(p, t) \leq \min_{p \in \bullet t} \max'_p = \nu(t)$.

As no urgency constraint is ever violated (the timed marking and configurations are reachable), we also have $\nu(t) = \min_{p \in \bullet t}(\max_p) \leq \alpha(t) + U(t) = A(t) + B(t) - A(t) = B(t)$. Hence t is fireable from (m, ν) , obtaining marking m_2 with $m_2(p) = m'_2(p)$ for all $p \in P$. The proof that (P1) holds for $(m_2, \nu_2), m'_2$ is exactly the same as above. As a conclusion, we have that \mathcal{R} is a bisimulation. \square

Proof for Section 4

Proposition 2. *Control State reachability, Reachability and (place) boundedness are undecidable for Timed PNU.*

Proof. As mentioned earlier, the proof follows on the same lines as the undecidability proof in [JLL77]. The proof is a standard encoding of reachability for a 2 counters machine with zero test (Minsky Machine). A Minsky machine is given as a set of counters $\{C_1, \dots, C_k\}$, and a sequence of instructions $\text{inst}_1.\text{inst}_2 \dots \text{inst}_n$. Each Instruction inst_i is of the form $\text{inst}_i : \text{inc}(j, k)$, meaning that the machine increments counter C_j and moves to instruction inst_k , or $\text{inst}_i : \text{Jzdec}(j, k, k')$ meaning that the machines tests whether counter C_j is zero or whether it is strictly greater. If it is 0, the machine moves to counter k' . Else it decrements C_j and moves to instruction k . A well known result is that it is undecidable whether a Minsky machine with two counters can reach an instruction x .

We can now show how to encode any two-counter Minsky machine M with a Timed PNU $\mathcal{N}_M = (P_M, T_M, \bullet(\cdot)_M, (\cdot)_M, m_0, \gamma_M, U_M)$. The set of places P_M contains two places p_1, p_2 (one per counter), and one place q_i per instruction inst_i of the machine. The set of transitions T_M contains one transition t_i for each increment instruction $\text{inst}_i : \text{inc}(j, k)$, with $\bullet t_{iM} = \{q_i\}$ and $t_{iM} \bullet = \{p_j, q_k\}$. Further, $U_M(t_i) = 0$ and $\gamma(q_i, t_i) = [0, \infty)$.

T_M also contains a pair of transitions t_i, t_i^z for each decrement instruction $\text{inst}_i : \text{Jzdec}(j, k, k')$. Intuitively, t_i will simulate the move of the machine when counter C_j is not empty, and t_i^z the move of the machine when counter C_j is empty. We have $\bullet t_{iM} = \{q_i, p_j\}$ and $t_{iM} \bullet = \{q_k\}$. Further, $U_M(t_i) = 0$ and $\gamma_M(q_i, t_i) = \gamma_M(p_j, t_i) = [0, \infty)$, that is transition t_i fires as soon as both places p_i and q_j contain at least one token. For transition t_i^z , we have $\bullet t_{iM}^z = \{q_i\}$ and $t_{iM}^z \bullet = \{q_{k'}\}$. Further, $U(t_i^z) = 0$ and $\gamma_M(q_i, t_i^z) = [1, \infty)$, that is transition t_i^z fires exactly 1 time unit after place q_i is filled.

Let us show that urgency prevents firing of transition t_i^z from a marking m with $m^{\sharp}(q_i) = 1$ and $m^{\sharp}(p_j) > 0$. Let us suppose that the machine moves from a marking m^- to m , with place q_i marked. This means that place q_i contains a token with age 0. Notice that at a given instant, only one place $q_j, j \in \{1, \dots, n\}$ can be marked in

a reachable marking of \mathcal{N}_M . Suppose that transition t_i^z can fire from a marking m' with $m'^{\#}(q_i) = 1$ and $m'^{\#}(p_j) > 0$, obtained from m by letting at least one time unit elapse. This means that $m'(p_j) \geq 1$, and hence, at least one time unit has elapsed from marking m . However, transition t_i is urgent from any marking with at least one token in q_i and one token in p_j , and in particular from m . Thus time cannot elapse from m , a contradiction. So, our encoding of M is faithful and a run $t_{i_1} \cdots t_{i_n}$ can happen in \mathcal{N}_M iff run $inst_{i_1} \cdots inst_{i_n}$ can happen in M .

Hence a two-counter machine M reaches instruction x iff place q_x is reachable in \mathcal{N}_M . This reachability question can be equivalently encoded as a coverability problem where the marking to cover is $m(p_x) = \{0\}$ and $m(p) = \emptyset$ for all other places. It is also known that checking boundedness of a counter in a Minsky machine is undecidable. Hence, boundedness and place boundedness are also undecidable in general for Timed PNU. \square

Proofs and details for Section 5 and 6

Proof of Theorem 1

A first step in our proofs is to show that the semantics of bounded Timed PN with Urgency can be encoded in a bisimilar way by timed automata. Let X be a finite set of real-valued variables called clocks. We write $\mathcal{C}(X)$ for the set of *constraints* over X , $\mathcal{C}_{ub}(X)$ of *upper bound* constraints over X as usual [AD94].

Definition 5 (Timed Automata (TA) [AD94]). A timed automaton \mathcal{A} over Σ_ε is a tuple (Q, q_0, X, δ, I) where Q is a finite set of locations, $q_0 \in Q$ is the initial location, X is a finite set of clocks, $I \in \mathcal{C}_{ub}(X)^Q$ assigns an invariant to each location and $\delta \subseteq Q \times \mathcal{C}(X) \times \Sigma_\varepsilon \times 2^X \times Q$ is a finite set of edges.

A *valuation* v is a mapping in $\mathbb{R}_{\geq 0}^X$. For $R \subseteq X$, the valuation $v[R]$ is the valuation v' such that $v'(x) = v(x)$ when $x \notin R$ and $v'(x) = 0$ otherwise. Finally, constraints of $\mathcal{C}(X)$ are interpreted over valuations: we write $v \models \gamma$ when constraint γ is satisfied by v . The semantics of a TA $\mathcal{A} = (Q, q_0, X, \delta, I)$ is the transition system $\llbracket \mathcal{A} \rrbracket = (S, s_0, \rightarrow)$ where $S = \{(q, v) \in Q \times (\mathbb{R}_{\geq 0})^X \mid v \models Inv(q)\}$, $s_0 = (\ell_0, \mathbf{0})$ and \rightarrow is defined by **delay moves:** $(\ell, v) \xrightarrow{d} (\ell, v+d)$ if $d \in \mathbb{R}_{\geq 0}$ and $v+d \models Inv(\ell)$; and **discrete moves:** $(\ell, v) \xrightarrow{a} (\ell', v')$ if there exists some $e = (\ell, \gamma, a, R, \ell') \in E$ s.t. $v \models \gamma$ and $v' = v[R]$. The (untimed) language of \mathcal{A} is defined as that of $\llbracket \mathcal{A} \rrbracket$ and is denoted by $\mathcal{L}(\mathcal{A})$.

The encoding of bounded Timed PN with Urgency into timed automata is not a real surprise, as similar bisimilar encodings have already been proposed for bounded TPNs. [DDSS07] propose to compute a marking timed automaton from TPN: the construction starts from the finite set of markings, and creates transitions of the form (m, g_t, t, m') for any pair of markings such that $m' = m - \bullet t + t \bullet$ with guard $g_t = \nu(t) \in [A(t), B(t)]$. We can then obtain the construction of finite timed automata from bounded TPNUs.

Proposition 2. *Let \mathcal{N} be a K -bounded TPNU. Then one can compute a finite timed automaton $\mathcal{A}_{\mathcal{N}}^B$ that is timed bisimilar to \mathcal{N} .*

Proof. We set $\mathcal{A}_{\mathcal{N}}^B = (Q, q_0, X, \delta, I)$, where Q is a set of states, δ a transition relation, $I : Q \rightarrow CST$ is an invariant relation attaching a constraint on clock values to each state, and X is a set of clocks.

The set of clocks X is computed as follows: $X = \{x_{p,i} \mid p \in P \wedge i \in 1..K\}$. Intuitively, $x_{p,i}$ will measure the age of the i^{th} token in place p . We denote by CST the set of all linear constraints that can be defined over X . Note that as we associate one clock per token in each place, as a place can be marked with less than K tokens, all clocks are not meaningful. We will compute inductively a set $Q \subseteq 2^X \times CST$. Each state $q = (A, cst)$ of the automaton will remember a set A of active clocks, and a constraint cst on the values of clocks. This construction is almost the state class construction of [Lime & Roux, DEDS, 06], with the slight difference that an active clock measures the age of a tokens instead of measuring the time since a transition was last newly enabled. To each state $q = (A, cst)$ and valuation $\nu_q : q\mathbb{R}$, we can associate a timed marking $m_{q,\nu}$ such that $m_{q,\nu}(p) = \{\nu(x_{p,i}) \mid x_{p,i} \in X\}$. Similarly, to each clock selection A , we can associate a marking $m^\sharp(A)(p) = |\{x_{p,k} \in A\}|$. We will say that A enables t if for every place p in $\bullet t$, $m^\sharp(A)(p) \geq 1$.

Let us now define the transition relation. We do not need to define timed moves, that will simply consist in letting time elapse from a state of the automaton. However, transitions of the automaton have to encode discrete moves of the net. As already pointed out, from a timed marking m , a discrete firing of a transition t can result in a finite number of new markings m'_1, \dots, m'_k , as the tokens to be consumed from each place can be chosen non-deterministically. Note however that as places are bounded by K , the number of markings that can be reached is bounded by $K^{|\bullet t|}$.

Let $q = (A, cst)$, and let t be a transition such that $\forall p \in \bullet t$ there exists at least one clock $x_{p,i}$ in A . We define $EXS(A, t)$ as the set of sets of clocks that contain exactly one clock of A per place in $\bullet t$. To fire, a transition needs to select a set of tokens that will be consumed, that is choose a element Y of $EXS(A, t)$. However, to be fireable, every clock $x_{p,i}$ in Y has to satisfy $x_{p,i} \in \gamma(p, t)$. For every $Y \in EXS(A, t)$, we define $g_{t,Y} = \bigwedge_{x_{p,i} \in Y} x_{p,i} \in \gamma(p, t)$ as the guard attached to transition t when choosing

set Y of token clocks. For a fixed origin state $q = (A, cst)$, a transition t and a chosen set of clocks Y , we will denote by $q' = Next(q, t, Y)$ the state reached after firing t from q , provided clocks in Y satisfy $g_{t,Y}$. State q' is of the form $q' = (A', cst')$, where $A' = A \setminus Y \cup R(A, t, Y)$, where $R(A, t, Y) = \{x_{p_1, k_1}, \dots, x_{p_q, k_q}\}$ is a set of clocks that contains exactly one clock per place in t^\bullet , and such that each k_j is the minimal index of clock attached to place p_j that does not appear in $A \setminus Y$ (this avoids introducing non-determinism when reusing clocks). This index exists, as \mathcal{N} is bounded. The constraint attached to the new state is computed as follows:

We add to cst conjunction of constraints $x'_{p,j} \leq x_{p',i}$ for each pair of clocks in $R, A \setminus Y$. Intuitively, this additional constraint means that the age of newly created tokens must be smaller or equal to the age of remaining tokens. We then eliminate variables $x_{p,i} \in Y$, using a variable elimination technique such as Fourier-Motzkin. We then rename variables $x'_{p,j}$ to $x_{p,j}$ to obtains cst' . One can note that applying arbitrary number of conjunction of inequations of the form $x \leq c, c \leq x$ where c is a constant and x a clock and elimination of variables, we can only reach a bounded number of constraints from cst_0 . Hence, Q is finite.

We can now define δ as follows $\delta = \{(q = (A, cst), g_{t,Y}, t, R(A, t, Y), q' = (A', cst')) \mid g_{t,Y} \wedge cst \text{ is satisfiable} \wedge t \text{ enabled by } A\}$. One can remark that a transition t can be enabled from several states, as two distinct clock selections may represent the same markings. The initial state of the automaton is $q_0 = (A_0, cst_0)$ where A_0 is the set $A_0 = \bigcup \{x_{p,1}, x_{p,k} \mid m_0^\#(p) = k\}$, and cst_0 is the set of constraints $cst_0 = \{x = y \mid x, y \in A_0\}$

The last ingredient of our timed automaton $\mathcal{A}_{\mathcal{N}}$ is the invariant attached to each state. Invariants must guarantee that clock values must not increase beyond urgency of a transition, i.e one can not reach a clock valuation ν such that for each place of the preset of some enabled transition t there exists a clock $x_{p,i}$ such that $\alpha(t) + U(t) < \nu(x_{p,i}) \leq \beta(t)$. □

Now, we turn to the proof of Theorem 1. The proof is performed in two steps. First, given a Timed PNrU \mathcal{N} , we construct a 1-bounded (labeled) Timed Petri Net \mathcal{N}_1 which is *timed equivalent* to the bounded part \mathcal{N}_B of \mathcal{N} (it is not possible to get bisimilarity because Timed Petri Nets do not use urgency). This is done by first converting the bounded part of the Timed PNrU into a timed automaton using Prop. 2, and then using the following result from [BHR08] (Theorem 7 and Corollary 1) to convert timed automata to Timed Petri Nets with equivalent (finite) language.

Lemma 6. [BHR08] *Let A be a timed automaton. Then we can construct a K' -bounded Timed Petri Net \mathcal{N}_1 such that \mathcal{N}_1 and A are timed equivalent.*

Now, the second step shows that the original Timed PNrU \mathcal{N} is timed equivalent to the net \mathcal{N}_2 formed by adding unbounded places of \mathcal{N} to the Timed Petri Net \mathcal{N}_1 .

Lemma 1. *\mathcal{N} and \mathcal{N}_2 are timed equivalent.*

Proof. First, let $(t^1, d^1) \cdots (t^n, d^n)$ be a finite word in the language of \mathcal{N} , where t^i are transitions of \mathcal{N} and $d^1 < \cdots < d^n$ are dates at which t^n occurs (that is, $d^i - d^{i-1}$ units of time occurred between t^{i-1} and t^i). Let m be some timed marking reached in \mathcal{N} after $(t^1, d^1) \cdots (t^n, d^n)$. First, as $(t^1, d^1) \cdots (t^n, d^n)$ is a timed word of \mathcal{N} , it is also a timed word of \mathcal{N}_B , the bounded restriction of \mathcal{N} , and thus of \mathcal{N}_1 . Let m_1 be a timed marking reached in \mathcal{N}_1 after $(t^1, d^1) \cdots (t^n, d^n)$. We now show by induction that the timed word $(t^1, d^1) \cdots (t^n, d^n)$ is a timed word of \mathcal{N}_2 , reaching a timed marking $m_2 = m_1 \cup m_u$ with m_u the restriction of m to P_u . First, it is trivial for the empty timed word.

Induction step: Let $\delta = d^n - d^{n-1}$ and $t = t^n$. Let m'_1, m' be two markings such that $m'_1 \xrightarrow{\delta} m'_1 + \delta \xrightarrow{t} m_1$ in \mathcal{N}_1 and $m' \xrightarrow{\delta} m' + \delta \xrightarrow{t} m$ in \mathcal{N} . By the induction hypothesis, the timed marking $m'_2 = m'_1 \cup m'_u$ is reached in \mathcal{N}_2 after $(t^1, d^1) \cdots (t^{n-1}, d^{n-1})$. We want to show that from m'_2 , δ timed units can elapse and then t is allowed, leading to m_2 . First, δ time units can elapse from m'_2 in \mathcal{N}_2 as there is no urgency in \mathcal{N}_2 (it is a Timed Petri Net). Further, tokens consumed in $m'_1 + \delta \xrightarrow{t} m_1$ can be consumed in $m'_2 + \delta$ as well, and so do tokens consumed from P_u in the transition $m' + \delta \xrightarrow{t} m$. Thus we indeed have $m'_2 \xrightarrow{\delta} m'_2 + \delta \xrightarrow{t} m_2$ and the proof is completed by induction.

Conversely, take $(t^1, d^1) \cdots (t^n, d^n)$ be a finite word in the language of \mathcal{N}_2 . Denote by $m_2 = m_1 \cup m_u$ a timed marking reached by this timed word, with m_1 a timed marking of \mathcal{N}_1 and m_u the timed marking over P_u the infinite places of \mathcal{N} . We have that $(t^1, d^1) \cdots (t^n, d^n)$ is a timed word of \mathcal{N}_1 and thus of \mathcal{N}_B the bounded part of \mathcal{N} . Let m_b be a timed marking of \mathcal{N}_B reached by this timed word. Then we prove by induction that $(t^1, d^1) \cdots (t^n, d^n)$ is a timed word of \mathcal{N} reaching $m = m_b \cup m_u$. The initial step is trivial for the empty timedword.

Induction step: Let $\delta = d^n - d^{n-1}$ and $t = t^n$. Let m'_2 such that $m'_2 \xrightarrow{\delta} m'_2 + \delta \xrightarrow{t} m_2$ in \mathcal{N}_2 , with $m'_2 = m'_1 \cup m'_u$ decomposed between P_1 and P_u , and let m'_b such that $m'_b \xrightarrow{\delta} m'_b + \delta \xrightarrow{t} m_b$ in \mathcal{N}_B the bounded part of \mathcal{N} . By the induction hypothesis, the timed marking $m' = m'_b \cup m'_u$ is reached in \mathcal{N}_2 after $(t^1, d^1) \cdots (t^{n-1}, d^{n-1})$. We want to show that from m' , δ timed units can elapse and then t is allowed, leading to m . First, δ time units can elapse from m'_b in \mathcal{N}_B , and thus it can elapse from m' in \mathcal{N} because \mathcal{N} is a Timed PNrU. Further, tokens consumed from $m'_b + \delta \xrightarrow{t} m_b$ can be consumed in m' as well, and so do tokens consumed from P_u in the transition $m'_2 + \delta \xrightarrow{t} m_2$. Thus we indeed have $m' \xrightarrow{\delta} m' + \delta \xrightarrow{t} m$ and the proof is completed by induction. \square

Thus, from Lemma 1 we obtain that for each Timed PNrU \mathcal{N} , one can construct a Timed Petri Net \mathcal{N}_2 such that \mathcal{N} and \mathcal{N}_2 are timed equivalent. By Theorem 2, we can then conclude the proof of Theorem 1.

Proof of Proposition 3 for the timed semantics

The exact proposition we prove for the timed semantics is the following.

Proposition 5. *For each TPN with restricted constraints \mathcal{N} , one can construct a Petri Net bisimilar to $\text{Timed}(\mathcal{N})$.*

The proof is again in two steps. First, given a TPNrU \mathcal{N} , we denote by \mathcal{N}_B its bounded part. We construct a 1-bounded (labeled) Petri Net \mathcal{N}_1 which is bisimilar to $\text{Timed}(\mathcal{N}_B)$. As $\text{Timed}(\mathcal{N}_B)$ is a bounded Timed PNU, we use proposition 2 to build a timed automaton \mathcal{A}_1 bisimilar to $\text{Timed}(\mathcal{N}_B)$. We interpret the regions automaton of this timed automata as a 1-safe Petri Net. Note that this net has no timed constraints. We call the resulting Petri Net $\mathcal{N}_1 = (P_1, T_1, \bullet, \bullet, m_{0,1})$. As \mathcal{A}_1 is time bisimilar to $\text{Timed}(\mathcal{N}_B)$, and \mathcal{N}_1 is (untimed) bisimilar to \mathcal{A}_1 , we have that \mathcal{N}_1 is (untimed) bisimilar to $\text{Timed}(\mathcal{N}_B)$.

Now, the second step shows that the original TPNrU $\text{Timed}(\mathcal{N})$ is bisimilar to the Petri net obtained by adding unbounded place of \mathcal{N} to \mathcal{N}_1 . Formally, we construct the Petri Net $\mathcal{N}_2 = (P_2, T_1, \bullet, \bullet, m_{0,2})$ as previously:

- The set P_2 of places of \mathcal{N}_2 is $P_2 = P_1 \cup P_u$, where P_u denotes the unbounded places of \mathcal{N} .
- Initial marking $m_{0,2}$ is the union of $m_{0,1}$ and of the restriction of the initial marking m_0 of \mathcal{N} to its unbounded place P_u .
- Finally, the set of transitions of \mathcal{N}_2 is the set T_1 of transitions of \mathcal{N}_1 , and the flow relations are defined in the following way. Let $t_1 \in T_1$ corresponding to a transition $t \in T$ in the original net \mathcal{N} . We have $p \in {}^*t_1$ 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 was an arc from p to t in \mathcal{N} .

Similarly, we have $p \in t_1^*$ if $p \in P_1$ and $p \in t_1 \bullet$ (arc from t_1 to p in \mathcal{N}_1), or if $p \in P_u$ and there is an arc from t to p in \mathcal{N} .

Lemma 6. *Timed(\mathcal{N}) and \mathcal{N}_2 are (untime) 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 into $m_2 = m_1 \cup m_u$ again by restriction to bounded/unbounded places. From proposition 5 and from the construction of \mathcal{N}_1 , we know that $\text{Timed}(\mathcal{N}_B)$ is bisimilar to \mathcal{N}_1 , and we denote by $R_{B,1}$ the unique largest bisimulation from timed markings of $\text{Timed}(\mathcal{N}_B)$ to markings of \mathcal{N}_1 .

We denote by R a relation from timed markings of $\text{Timed}(\mathcal{N})$ to markings of \mathcal{N}_2 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 $m'_u = m_u^\#$. Obviously, we have $(m_0, m_{0,2}) \in R$. 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, $m_u^\# \geq \bullet t \cap P_u$. 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 $m_1 \cup m_u^\# \xrightarrow{t} m'_1 \cup m'_u^\#$. Furthermore $m'_u^\# = |m'_u$, so $(m', m'_1 \cup m'_u^\#) \in R$.

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, t as t can fire, we have $m_1 \xrightarrow{t} m'_1$. So, there exists a 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 $m_b \xrightarrow{\delta} m_b + \delta \xrightarrow{t} m'_b$.

Now, \mathcal{N} is a TPN with restricted constraints. First, it means that $m \xrightarrow{\delta} m + \delta$ does not violate any urgency constraints, as they are all in the bounded part of \mathcal{N} . Now, 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 $(m'_u)^\# = m'_3$. First, because $(m, m_2) \in R$, with $m_2 = m_1 \cup m_3$, we have $m^\#(p) = m_3(p) \geq 1$ for all $p \in P_u \cap \bullet t$. Also, we have trivially that $m^\#(p) = m_b^\#(p) \geq 1$ 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 $(m'_u)^\# = m'_3$. For the bounded part, we chose to consume the tokens consumed during the transition $m_b + \delta \xrightarrow{m'_b} m'_b$. We thus obtain $m' = m'_b + m'_u$, and $(m', m'_2) \in R$. Hence R is a bisimulation relation. \square

Lemma 6 thus complete the proof of proposition 5 for the timed semantics. Thus we can decide reachability in $\text{Timed}(\mathcal{N})$ for \mathcal{N} a TPN with restricted constraints.

Proof of Theorem 4

Let (\mathcal{N}, P, Q) , with $\mathcal{N} = (P \cup Q, T, \bullet(), ()^\bullet, m_0, \lambda, I)$ be a TPNrU, where P is the set of bounded states and Q the set of unbounded states. We prove the two remaining lemmas.

Lemma 5. *Let m be a reachable marking of $\text{Reach}(\text{Timed}(\mathcal{N}))$. Then one can reach in $\text{Timed}(\mathcal{N}')$ any m' with:*

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

Proof. of Lemma 5. We proceed by induction on the length of run reaching m in \mathcal{N} . For run of length 0, this is trivial. Assume that m is reached after applying t from m^- . Let m' be any marking satisfying the condition above wrt m . Let us show that we can reach m' in $\text{Timed}(\mathcal{N}')$.

Assume that t let time elapse by $\delta > 0$ units of time. For all $p \in P \cup Q$, for all $age_p \in m(p)$, $age_p \geq \delta$. We first show that for all $p' \in P'$, for all $age'_p \in m'(p')$, $age'_p \geq \delta$. For $q = 1_p^t$ for some t, p , we have $m'(q) = \{\min(m(p_t), age_p)\}$ with $age_p \in m(p)$. Now, we have $m(p_t) \geq \delta$ and $age_p \geq \delta$. For other $p' \in P'$, $age'_p \in m'(p)$, there exists $p \in P \cup Q$ with $age'_p \in m(p)$. We thus obtain for all $p' \in P'$, for all $age'_p \in m'(p')$, $age'_p \geq \delta$. Hence we can define $m'' = m' - \delta$ as a marking.

The number of tokens in m, m^- is the same, and the number of tokens in m', m'' is the same (only timing changed). Thus, as $m \equiv m'$, we have $m^- \equiv m''$. Now, $m^- \equiv m''$ satisfies the conditions above, we can thus apply the induction hypothesis telling us that m'' is reachable in $\text{Timed}(\mathcal{N}')$. Now, waiting δ units of time from m'' is allowed in $\text{Timed}(\mathcal{N}')$ as it does not violate any urgency. Indeed, by contradiction, if the urgency of t is violated, it implies that $I(t) = [a, b]$ and thus $\bullet t \subseteq P$ is made of only bounded places. As m^- and m'' coincide on bounded places, it would also violate urgency on m^- , which is by definition not the case, a contradiction. Hence δ units of time can elapse from m'' , and the marking reached is m' . That is, m' is reachable in $\text{Timed}(\mathcal{N}')$.

Now, consider t a discrete transition. Let say $(age_p)_{p \in \bullet t}$ are the tokens which are consumed by t from m^- to m . Let $S_q^- = \{0\}$ if $q \in t^\bullet$ and \emptyset otherwise. We have $m(p_t) = 0$. Thus for all $q \in \bullet t$, $m'(1_q^t) = \emptyset$ or $m'(1_q^t) = 0$.

We define m'' in the following way with $m^- \equiv m''$:

- for all $p \in P \cup P_T$, we have $m''(p) = m^-(p)$.
- For all $q \in Q$, let t_q be the (possibly ϵ) transition and $nage_q$ be the (possible ϵ) token age with
 - If $t \in T'_q$ (in particular $q \in \bullet t$ and $m'(1_q^t) = \{0\}$), then $t_q = t$, and $nage_q$ is the token associated with $m'(1_q^t) = \{\min(nage_q, 0)\}$,
 - Else, $(m'(1_q^t) = \emptyset)$ if $q \in t^\bullet$ and $0 \notin m'(q)$, then we let t_q be any transition with $m'(1_q^t) = \{0\}$, and $nage_q = 0$,

- Else $t_q = nage_q = \epsilon$,
- In both case, a token of $nage_q \in m'(q)$ has just been consumed with a $start_q^{t_q}$ to create a token with age 0 in $m'(1_q^t)$.
- for all $q \in Q$, we have $m''(q) = m'(q) \sqcup \{nage_q\} \setminus S_q^-$,
 - for all $t' \in T, q \in \bullet t'$,
 - if $t' = t$ then $m''(1_q^t) = \min(m^-(p_t), age_q)$,
 - else if $t' = t_q$ then $m''(1_q^t) = \emptyset$,
 - else $m''(1_q^t) = m'(1_q^t)$,

Let show that $m^- \equiv m''$ satisfies the hypothesis. First, for every $p \in P \cup P_t$, this is trivial. Now let $q \in Q$. We have $m^-(q) \sqcup S_q^- = m(q) \sqcup S_q^-$. We have $m(q) = \{age_{tt} \mid tt \in T'_q\} \sqcup m'(q)$.

Hence $m^-(q) = m'(q) \sqcup \{age_{tt} \mid tt \in T'_q \setminus \{t_q\}\} \sqcup \{age_{t_q}\} \sqcup S_q^- \setminus S_q^-$.

Now, for all $tt \in T'_q \setminus \{t_q\}$ (in particular, $tt \neq t$), we have $m''(1_q^t) = m'(1_q^t)$, thus this part is covered. Also, if $q \in \bullet t$, then $S_q = \{age_q\}$ is the token associated with $m''(1_q^t) = \min(m^-(p_t), age_q)$.

It remains $m'(q) \sqcup \{age_{t_q}\} \setminus S_q^-$ to associate with $m''(q) = m'(q) \sqcup \{nage_q\} \setminus S_q^-$, which is trivial (either $nage_q = 0$ and any token can be associated with it, or else $nage_q$ is by definition the token associated with age_{t_q}).

Hence by the induction hypothesis, m'' is reachable in $\text{Timed}(\mathcal{N}')$. Now performing t from m'' then $start_q^{t_q}$ for all q such that $t_q \neq \epsilon$, we obtain m' . That is, m' is reachable in $\text{Timed}(\mathcal{N}')$. \square

We turn now to the proof of the other lemma.

Lemma 4. *Let m' be a reachable marking of $\text{Reach}(\text{Timed}(\mathcal{N}'))$. Then there exists a marking $m \in \text{Reach}(\text{Timed}(\mathcal{N}))$ with $m \equiv m'$.*

Proof. of Lemma 4.

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

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

For $m' = m'_0$, this is trivial. Induction step: Let m' be a reachable marking of $\text{Timed}(\mathcal{N})$. A path reaching m' ends with $m'' \mapsto (t)m'$. Hence m'' is reached in less steps than m' , and we can apply the induction hypothesis. One can reach m in $\text{Timed}(\mathcal{N}')$ with:

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

Case: the transition is $start_p^t$. Then it is easy to see that $m'(q) = m(q)$ for all q but for $m(p) = m'(p) \sqcup \{age\}$, $m(1_q^t) = \emptyset$ and $m'(1_q^t) = 0$. It is easy to check that m satisfies the hypothesis wrt m' . As m is known to be reachable in $Timed(\mathcal{N})$, we conclude.

Case t makes time elapse by δ unit of time. m and m'' agrees on the bounded places, hence as the urgency is not violated in m'' by elapsing δ unit of time, it is not violated in m by elapsing δ unit of time. Thus $m + \delta$ is reachable in $Timed(\mathcal{N})$. Last, it is easy to see that $m + \delta$ satisfies the hypothesis wrt $m' = m'' + \delta$.

Case $t \in T$. If $\bullet t$ has only bounded places, then m'' and m agree on these places, and one can apply t from m to obtain m^+ which satisfies the hypothesis wrt m' . Else, consider an unbounded place $q \in \bullet t \cap Q$. Thus $I(t) = [a, +\infty)$. We have $m'(1_q) = \{nage_q\}$ with $nage_q \geq a$. Thus there exists $age'_q \in m(q)$ with $age'_q \geq nage_q$. In particular, t is enabled wrt q . This is true for all q , so t is enabled from m . Denote m^+ the reached configuration of $Timed(\mathcal{N})$ after deleting the chosen tokens age_q with $age_q \geq nage_q$.

Now, firing t in m creates the same tokens in the same place with the same age 0 as firing t from m' . It deletes token $m'(1_q)$ and one token in q of age bigger or equal. Hence we still have for all $q \in Q$, $m^+(q) = m'(q) \sqcup \{age_1, \dots, age_k\}$ and there exists a bijection $f : T_q \mapsto [1, k]$ with $m'(1_q^t) \leq age_{f(t)}$ for all $t \in T_q$. That is, m^+ satisfies the hypothesis wrt m' . \square