

Integrated Modeling and Verification of Real-Time Systems through Multiple Paradigms

Marcello M. Bersani, Carlo A. Furia, Matteo Pradella, and Matteo Rossi

July 2009

Abstract

Complex systems typically have many different parts and facets, with different characteristics. In a multi-paradigm approach to modeling, formalisms with different natures are used in combination to describe complementary parts and aspects of the system. This can have a beneficial impact on the modeling activity, as different paradigms can be better suited to describe different aspects of the system. While each paradigm provides a different view on the many facets of the system, it is of paramount importance that a coherent comprehensive model emerges from the combination of the various partial descriptions. In this paper we present a technique to model different aspects of the same system with different formalisms, while keeping the various models tightly integrated with one another. In addition, our approach leverages the flexibility provided by a bounded satisfiability checker to encode the verification problem of the integrated model in the propositional satisfiability (SAT) problem; this allows users to carry out formal verification activities both on the whole model and on parts thereof. The effectiveness of the approach is illustrated through the example of a monitoring system.

Keywords: Metric temporal logic, timed Petri nets, timed automata, discretization, dense time, bounded model checking.

1 Introduction

Modeling paradigms come in many different flavors: graphical or textual; executable or not; formal, informal, or semi-formal; more or less abstract; with different levels of expressiveness, naturalness, conciseness, etc. Notations for the design of real-time systems, in addition, include a notion of time, whose characteristics add a further element of differentiation [14].

A common broad categorization of modeling notations distinguishes between *operational* and *descriptive* paradigms [10]. Operational notations — such as Statecharts, finite state automata, or Petri nets — represent systems through the notions of *state* and *transition* (or event); system behavior consists in evolutions from state to state, triggered by event occurrences. On the other hand, descriptive paradigms — such as temporal logics, descriptive logics, or algebraic formalisms — model systems by declaring their fundamental *properties*.

The distinction between operational and descriptive models is, like with most classifications, neither rigid nor sharp. Nonetheless, it is often useful in practice to guide the developer in the choice of notation based on what is being modeled and what are the ultimate goals (and requirements) of the modeling endeavor. In fact, operational

and descriptive notations have different — and often complementary — strengths and weaknesses. Operational models, for instance, are often easier to understand by experts of domains other than computer science (mechanical engineers, control engineers, etc.), which makes them a good design vehicle in the development of complex systems involving components of many different natures. Also, once an operational model has been built, it is typically straightforward to execute, simulate, animate, or test it. On the other hand, descriptive notations are the most natural choice when writing partial models of systems, because one can build the description *incrementally* by listing the (partial) known properties one at a time. For similar reasons, descriptive models are often excellent languages to document the *requirements* of a system: the requirements elicitation process is usually an incremental trial-and-error activity, and thus it benefits greatly from notations which allow cumulative development.

When modeling timed systems, in addition, the choice of the time domain is a crucial one, and it can significantly impact on the features of the model [10]. For example, a dense time model is typically needed to represent true asynchrony. Discrete time, instead, is usually more amenable to automated verification, and is at the basis of a number of quite mature techniques and tools that can be deployed in practice to verify systems.

In this paper we present a technique to model different aspects of the same system with different formalisms, while keeping the various models tightly integrated with one another. In this approach, modelers can pick their preferred modeling technique and modeling paradigm (e.g., operational or descriptive, continuous or discrete) depending on the particular facet or component of the system to be described. Integration of the separate snippets in a unique model is made possible by providing a common formal semantics to the different formalisms involved. Finally, our approach leverages the flexibility provided by a bounded satisfiability checker to encode the verification problem of the integrated model in the propositional satisfiability (SAT) problem; this allows users to carry out formal verification activities both on the whole model and on parts thereof.

The technique presented in this paper hinges on Metric Temporal Logic (MTL) to provide a common semantic foundation to the integrated formalisms, and on the results presented in [13] to integrate continuous- and discrete-time MTL fragments into a unique formal description. Operational formalisms can then be introduced in the framework by providing suitable MTL formalizations, which can then be discretized as well according to the same technique. While this idea is straightforward in principle, putting it into practice is challenging for several basic reasons. First, in order to have full discrete-time decidability we have to limit ourselves to *propositional* MTL [4]; its relatively limited expressive power makes it arduous to formalize completely the behavior of operational models (some technical facts, briefly described in Section 2, justify this intuition). Second, even if we used a more expressive first-order temporal-logic language, formalizing the semantics of “graphical” operational formalisms is usually tricky as several semantic subtleties that are “implicit” in the original model must be properly understood and resolved when translating them into a logic language. See for instance extensive discussions of such subtleties in [8] for timed Petri nets and in [19] for Statecharts. Third, not any MTL axiomatization is amenable to the discretization techniques of [11], as syntactically different MTL descriptions yielding the same underlying semantics provide discretizations of wildly different “qualities”. Indeed, experience showed that the most “natural” axiomatizations of operational formalisms require substantial rewriting in order to work reasonably well under the discretization framework. Crafting suitable MTL descriptions has proved demanding, delicate,

and crucially dependent on the features of the operational formalism at hand. In this respect, our previous work [12] focused on a variant of Timed Automata (TA) — a typical “synchronous” operational formalism. The formalization of intrinsically asynchronous components — such as those that sit at the boundary between the system and its environment — demands however the availability of a formalism that is both operational and “asynchronous”. To this end, the present paper develops an axiomatization of Timed Petri Nets (TPN), an “asynchronous” operational formalism, integrates all three formalisms (MTL, TA, and TPN) into a unique framework, and evaluates an implementation of the framework on a monitoring system example.

The paper is structured as follows. Section 1.1 briefly discusses some works that are related to the approach and technique presented in this article. Section 2 introduces the relevant results on which the modeling and verification approach presented in this paper are based; more precisely, the section introduces MTL, timed automata and their MTL-based semantics, and the discretization technique for continuous-time MTL formulas. Section 3 presents the (continuous-time) MTL semantics of timed Petri nets and uses it to derive a discretized version of timed Petri nets that can be input to verification engines for discrete-time MTL (e.g., *Zot*). Section 4 shows how the various formalisms can be used to describe, and then combine together in a unique model, different aspects and parts of the same system; in addition, it reports on some verification tests carried out on the modeled system. Finally, Section 5 concludes and outlines some future works in this line of research.

1.1 Related work

Combining different modeling paradigms in a single framework for verification purposes is not a novel concept. In fact, there is a rich literature on dual-language approaches, which combine an operational formalism and a descriptive formalism into one analysis framework [10]. The operational notation is used to describe the system dynamics, whereas the properties to be checked are expressed through the descriptive notation. Model-checking techniques [7] are a widely-used example of a dual-language approach to formal verification. Dual-language frameworks, however, usually adopt a rigid stance, in that one formalism is used to describe the system, while another is used for the properties to be verified. In this work we propose a flexible framework in which different paradigms can be mixed for different design purposes: system modeling, property specification and also verification.

Modeling using different paradigms is a staple of UML [18]. In fact, the UML modeling language is actually a blend of different notations (message sequence charts, Statecharts, OCL formulas, etc.) with different characteristics. The UML framework provides means to describe the same (software) systems from different, possibly complementary, perspectives. However, the standard language is devoid of mechanisms to guarantee that an *integrated* global view emerges from the various documents or that, in other words, the union of the different views yields a precise, coherent model.

Some work has been devoted to the (structural) transformation between models to re-use verification techniques for different paradigms and to achieve a unified semantics, similarly to the approach of this paper. Cassez and Roux [5] provide a structural translation of TPN into TA that allows one to piggy-back the efficient model-checking tools for TA. Our approach is complementary to [5] and similar works¹ in several ways. First, our transformations are targeted to a discretization framework: on the one hand,

¹See the related work section of [5] for more examples of transformational approaches.

this allows a more lightweight verification process as well as the inclusion of discrete-time components within the global model; on the other hand, discretization introduces incompleteness that might reduce its effectiveness. Second, we leverage on a descriptive notation (MTL) rather than an operational one. This allows the seamless integration of operational and descriptive components, whereas the transformation of [5] stays within the model-checking paradigm where the system is modeled within the operational domain and the verified properties are modeled with a descriptive notation. Also, state-of-the-art of tools for model-checking of TA (and formalisms of similar expressive power) do not support full real-time temporal logics (such as TCTL) but only a subset of significantly reduced expressive power. We claim that the model and properties we consider in the example of Section 4 are rather sophisticated and deep—even after weighting in the inherent limitations of our verification technique.

For the sake of brevity, we omit in this report a description of related works on the discretization of continuous-time models. The interested reader can refer to [11] for a discussion of this topic.

2 Background

2.1 Continuous- and discrete-time real-time behaviors

We represent the concept of *trace* (or *run*) of some real-time system through the notion of *behavior*. Given a time domain \mathbb{T} and a finite set \mathcal{P} of atomic propositions, a behavior b is a mapping $b : \mathbb{T} \rightarrow 2^{\mathcal{P}}$ which associates with every time instant $t \in \mathbb{T}$ the set $b(t)$ of propositions that hold at t . $\mathcal{B}_{\mathbb{T}}$ denotes the set of all behaviors over \mathbb{T} (for an implicit fixed set of propositions). $t \in \mathbb{T}$ is a *transition point* for behavior b iff t is a discontinuity point of the mapping b . Depending on whether \mathbb{T} is a discrete, dense, or continuous set, we call a behavior over \mathbb{T} discrete-, dense-, or continuous-time respectively. In this report, we assume the natural numbers \mathbb{N} as discrete time domain and the nonnegative real numbers $\mathbb{R}_{\geq 0}$ as continuous (and dense) time domain.

Non-Zeno and non-Berkeley. Over continuous-time domains, it is customary to consider only physically meaningful behaviors, namely those respecting the so-called non-Zeno property. A continuous-time behavior b is non-Zeno if the sequence of transition points of b has no accumulation points. For a non-Zeno behavior b , it is well-defined the notions of values to the left and to the right of any transition point $t > 0$, which we denote as $b^-(t)$ and $b^+(t)$, respectively. When a proposition $p \in \mathcal{P}$ is such that $p \in b^-(t) \Leftrightarrow p \notin b^+(t)$ (i.e., p switches its truth value about t), we say that p is “triggered” at t . In order to ensure reducibility between continuous and discrete time, we consider non-Zeno behaviors with a stronger constraint, called *non-Berkeleyness*. A continuous-time behavior b is non-Berkeley for some positive constant $\delta \in \mathbb{R}_{>0}$ if, for all $t \in \mathbb{T}$, there exists a closed interval $[u, u + \delta]$ of size δ such that $t \in [u, u + \delta]$ and b is constant throughout $[u, u + \delta]$. Notice that a non-Berkeley behavior (for any δ) is non-Zeno *a fortiori*. The set of all non-Berkeley continuous-time behaviors for $\delta > 0$ is denoted by $\mathcal{B}_{\chi}^{\delta} \subset \mathcal{B}_{\mathbb{R}_{\geq 0}}$. In the following we always assume behaviors to be non-Berkeley, unless explicitly stated otherwise.

Syntax and semantics. From a purely semantic point of view, one can consider the model of a (real-time) system simply as a set of behaviors [3, 9] over some time domain \mathbb{T} and sets of propositions. In practice, however, systems are modeled through some

suitable notation: in this paper we consider a mixture of MTL formulas [15, 4], TA [1, 2], and TPN [6]. Given an MTL formula, a TA, or a TPN μ , and a behavior b , $b \models \mu$ denotes that b represents a system evolution which satisfies all the constraints imposed by μ . If $b \models \mu$ for some $b \in \mathcal{B}_{\mathbb{T}}$, μ is called \mathbb{T} -satisfiable; if $b \models \mu$ for all $b \in \mathcal{B}_{\mathbb{T}}$, μ is called \mathbb{T} -valid. Similarly, if $b \models \mu$ for some $b \in \mathcal{B}_{\chi}^{\delta}$, μ is called χ^{δ} -satisfiable; if $b \models \mu$ for all $b \in \mathcal{B}_{\chi}^{\delta}$, μ is called χ^{δ} -valid.

2.2 Descriptive notation: Metric Temporal Logic

Let \mathcal{P} be a finite (non-empty) set of atomic propositions and \mathcal{I} be the set of all (possibly unbounded) intervals of the time domain \mathbb{T} with rational endpoints. We abbreviate intervals with pseudo-arithmetic expressions, such as $= d$, $< d$, $\geq d$, for $[d, d]$, $(0, d)$, and $[d, +\infty)$, respectively.

MTL syntax. The following grammar defines the syntax of (propositional) MTL, where $I \in \mathcal{I}$ and $p \in \mathcal{P}$.

$$\phi ::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid U_I(\phi_1, \phi_2) \mid S_I(\phi_1, \phi_2)$$

The basic temporal operators of MTL is the *bounded until* $U_I(\phi_1, \phi_2)$ (and its past counterpart *bounded since* S_I) which says that ϕ_1 holds until ϕ_2 holds, with the additional constraint that ϕ_2 must hold within interval I . Throughout the paper we omit the explicit treatment of past operators (i.e., S_I and derived) as it can be trivially derived from that of the corresponding future operators.

MTL semantics. MTL semantics is defined over behaviors, parametrically with respect to the choice of the time domain \mathbb{T} . While the semantics of Boolean connectives and In particular, the definition of the *until* operators is as follows:

$$\begin{aligned} b(t) \models_{\mathbb{T}} U_I(\phi_1, \phi_2) & \quad \text{iff} \quad \text{there exists } d \in I \text{ such that: } b(t+d) \models_{\mathbb{T}} \phi_2 \\ & \quad \text{and, for all } u \in [0, d] \text{ it is } b(t+u) \models_{\mathbb{T}} \phi_1 \\ b \models_{\mathbb{T}} \phi & \quad \text{iff} \quad \text{for all } t \in \mathbb{T}: b(t) \models_{\mathbb{T}} \phi \end{aligned}$$

We remark that a global satisfiability semantics is assumed, i.e., the satisfiability of formulas is implicitly evaluated over *all* time instants in the time domain. This permits the direct and natural expression of most common real-time specifications (e.g., time-bounded response, time-bounded invariance, etc.) without resorting to nesting of temporal operators.

Granularity. For an MTL formula ϕ , let \mathcal{J}_{ϕ} be the set of all non-null, finite interval bounds appearing in ϕ . Then, \mathcal{D}_{ϕ} is the set of positive values δ such that any interval bound in \mathcal{J}_{ϕ} is an integer if divided by δ .

2.2.1 Derived (temporal) operators.

It is customary to introduce a number of derived (temporal) operators, to be used as shorthands in writing specification formulas. We assume a number of standard abbreviations such as \perp , \top , \vee , \Rightarrow , \Leftrightarrow ; when $I = (0, \infty)$, we drop the subscript interval in temporal operators. All other derived operators used in this paper are listed in Table 1 ($\delta \in \mathbb{R}_{>0}$ is a parameter used in the discretization techniques, discussed shortly). In the following we describe briefly and informally the purpose of such derived operators, focusing on future ones (the meaning of the corresponding past operators is easily derivable).

- For propositions in the set $\{\gamma(x) \mid x \in X\}$, $\odot_{x \in X \subseteq Y} \gamma(x)$ states that $\gamma(x)$ holds for all x in X and does not hold for all x in the complement set $Y \setminus X$.
- A few common derived temporal operators such as $R_I, \diamond_I, \square_I$ are defined with the usual meaning: R_I (*release*) is the dual of the *until* operator; $\diamond_I(\phi)$ means that ϕ happens within time interval I in the future; $\square_I(\phi)$ means that ϕ holds throughout the whole interval I in the future.
- $\widetilde{\circ}(\phi)$ and $\circ(\phi)$ are useful over continuous time only, and describe ϕ holding throughout some unspecified non-empty interval in the strict future; more precisely, if t is the current instant, there exists some $t' > t$ such that ϕ holds over $\langle t, t' \rangle$, where the interval is left-open for $\widetilde{\circ}$ and left-closed for \circ .
- \triangle and \blacktriangle describe different types of *transitions*. Namely, $\triangle(\phi_1, \phi_2)$ describes a switch from ϕ_1 to ϕ_2 , irrespective of which value holds at the current instant, whereas $\blacktriangle(\phi_1, \phi_2)$ describes a switch from ϕ_1 to ϕ_2 such that ϕ_1 holds at the current instant and ϕ_2 will hold in the immediate future. Note that if $\triangle(\phi_1, \phi_2)$ holds at some instant t , $\blacktriangle(\phi_1, \phi_2)$ holds over $(t - \delta, t)$.
- $\triangle(\phi), \blacktriangle(\phi)$ are shorthands for transitions of a single item; correspondingly the \wr, \wr, \mathfrak{W} “trigger” operators are introduced: $\wr(\phi)$ denotes a transition of ϕ from false to true or *vice versa*, whereas $\mathfrak{W}(\phi)$ describes a similar transition where the value of ϕ at the current instant is unspecified. $\mathfrak{W}(\phi' \rightsquigarrow \phi)$ describes a more complex transition of ϕ , one which is “triggered” by the auxiliary proposition ϕ' .
- It is also convenient to introduce the “dual” operators $\bar{\wr}, \bar{\mathfrak{W}}$ which describe “non-transitions” of their argument. For instance, $\bar{\wr}(\phi)$ says that the truth value of ϕ (whatever it is) does not change from the current instant to the immediate future.
- Finally, $\text{Alw}(\phi)$ expresses the invariance of ϕ . Since $b \models_{\mathbb{T}} \text{Alw}(\phi)$ iff $b \models_{\mathbb{T}} \phi$, for any behavior b , $\text{Alw}(\phi)$ can be expressed without nesting if ϕ is flat, through the global satisfiability semantics introduced beforehand.

2.3 Operational notations: Timed Automata and Timed Petri Nets

For lack of space, we omit a formal presentation of TA, which have been however introduced in the framework in previous work [12] and focus on MTL and TPN in the following. Section 4 will however informally illustrate the syntax and semantics of TA on an example, with a level of detail sufficient to understand its role within the framework.

Timed Petri nets syntax. A *Timed Petri Net* (TPN) is a tuple $N = \langle P, T, F, M_0, \alpha, \beta \rangle$:

- P is a finite set of *places*;
- T is a finite set of *transitions*;
- $F \subseteq (P \times T) \cup (T \times P)$ is the *flow relation*;
- $M_0 : P \rightarrow \mathbb{N}$ is the *initial marking*;
- $\alpha : T \rightarrow \mathbb{Q}_{\geq 0}$ gives the *earliest firing times* of transitions; and

OPERATOR	≡	DEFINITION
$\bigcirc_{x \in X \subset Y} \gamma(x)$	≡	$\bigwedge_{x \in X} \gamma(x) \wedge \bigwedge_{y \in Y \setminus X} \neg \gamma(y)$
$R_I(\phi_1, \phi_2)$	≡	$\neg U_I(\neg \phi_1, \neg \phi_2)$
$T_I(\phi_1, \phi_2)$	≡	$\neg S_I(\neg \phi_1, \neg \phi_2)$
$\underset{\leftarrow}{\diamond}_I(\phi)$	≡	$U_I(\top, \phi)$
$\underset{\rightarrow}{\diamond}_I(\phi)$	≡	$S_I(\top, \phi)$
$\square_I(\phi)$	≡	$R_I(\perp, \phi)$
$\overline{\square}_I(\phi)$	≡	$T_I(\perp, \phi)$
$\widetilde{\bigcirc}(\phi)$	≡	$U_{(0,+\infty)}(\phi, \top) \vee (\neg \phi \wedge R_{(0,+\infty)}(\phi, \perp))$
$\widetilde{\bigcirc}(\phi)$	≡	$S_{(0,+\infty)}(\phi, \top) \vee (\neg \phi \wedge T_{(0,+\infty)}(\phi, \perp))$
$\bigcirc(\phi)$	≡	$\phi \wedge \widetilde{\bigcirc}(\phi)$
$\overline{\bigcirc}(\phi)$	≡	$\phi \wedge \widetilde{\overline{\bigcirc}}(\phi)$
$\Delta(\phi_1, \phi_2)$	≡	$\begin{cases} \widetilde{\bigcirc}(\phi_1) \wedge (\phi_2 \vee \widetilde{\bigcirc}(\phi_2)) & \text{if } \mathbf{T} = \mathbf{R}_{\geq 0} \\ \widetilde{\diamond}_{=1}(\phi_1) \wedge \diamond_{[0,1]}(\phi_2) & \text{if } \mathbf{T} = \mathbf{N} \end{cases}$
$\blacktriangle(\phi_1, \phi_2)$	≡	$\begin{cases} \phi_1 \wedge \diamond_{=\delta}(\phi_2) & \text{if } \mathbf{T} = \mathbf{R}_{\geq 0} \\ \phi_1 \wedge \diamond_{=1}(\phi_2) & \text{if } \mathbf{T} = \mathbf{N} \end{cases}$
$\Delta(\phi)$	≡	$\Delta(\neg \phi, \phi)$
$\blacktriangle(\phi)$	≡	$\blacktriangle(\neg \phi, \phi)$
$\beth(\phi)$	≡	$\blacktriangle(\phi) \vee \blacktriangle(\neg \phi)$
$\beth(\phi)$	≡	$\Delta(\phi) \vee \Delta(\neg \phi)$
$\beth(\phi' \rightsquigarrow \phi)$	≡	$\begin{cases} \widetilde{\bigcirc}(\neg \phi) \wedge \square_{=\delta}(\phi' \Rightarrow \phi) \vee \widetilde{\bigcirc}(\phi) \wedge \square_{=\delta}(\phi' \Rightarrow \neg \phi) & \text{if } \mathbf{T} = \mathbf{R}_{\geq 0} \\ \widetilde{\square}_{[0,1]}(\neg \phi) \wedge \square_{[0,2]}(\phi' \Rightarrow \phi) \vee \widetilde{\square}_{[0,1]}(\phi) \wedge \square_{[0,2]}(\phi' \Rightarrow \neg \phi) & \text{if } \mathbf{T} = \mathbf{N} \end{cases}$
$\beth(\phi)$	≡	$\Delta(\phi, \phi)$
$\beth(\phi)$	≡	$\blacktriangle(\phi, \phi) \vee \blacktriangle(\neg \phi, \neg \phi)$
$\beth(\phi' \rightsquigarrow \phi)$	≡	$\begin{cases} \widetilde{\bigcirc}(\phi) \wedge \square_{=\delta}(\phi' \Rightarrow \phi) \vee \widetilde{\bigcirc}(\neg \phi) \wedge \square_{=\delta}(\phi' \Rightarrow \neg \phi) & \text{if } \mathbf{T} = \mathbf{R}_{\geq 0} \\ \widetilde{\square}_{[0,1]}(\phi) \wedge \square_{[0,2]}(\phi' \Rightarrow \phi) \vee \widetilde{\square}_{[0,1]}(\neg \phi) \wedge \square_{[0,2]}(\phi' \Rightarrow \neg \phi) & \text{if } \mathbf{T} = \mathbf{N} \end{cases}$
$\text{Alw}(\phi)$	≡	$\phi \wedge \square_{(0,+\infty)}(\phi) \wedge \overline{\square}_{(0,+\infty)}(\phi)$

Table 1: MTL derived temporal operators

- $\beta : T \rightarrow \mathbb{Q}_{\geq 0} \cup \{\infty\}$ gives the *latest firing times* of transitions.

In general, a mapping $M : P \rightarrow \mathbb{N}$ is called a *marking* of N . Given $a \in P \cup T$, let $\bullet a = \{b \mid bFa\}$ and $a\bullet = \{b \mid aFb\}$ denote the preset and postset of a , respectively. We assume that every node $a \in P \cup T$ has a nonempty preset or a nonempty postset (or both); this is clearly without loss of generality.

Timed Petri nets semantics. The semantics of TPN is usually given as sequences of transition firings and place markings; see [6] for formal definitions. Correspondingly, a TPN is called *k-safe* for $k \in \mathbb{N}$ iff for every reachable marking M it is $M(p) \leq k$ for all $p \in P$. A TPN that is *k-safe* for some $k \in \mathbb{N}$ is called *bounded*.

In this report we assume *1-safe TPN*. This allows a simplified description of the semantics, where any marking is completely described by a set $M \subseteq P$ of places such that a place is marked iff it is in M . We remark, however, that extending the presentation to generic *bounded TPN* would be routine. On the other hand, unbounded TPN would not be discretizable according to the notion of Section 2.4, hence they would fit only in a different framework. To further simplify the presentation, we assume non-Berkeley behaviors for some generic $\delta > 0$ in presenting the semantics; correspondingly we do not have to consider zero-time transitions as every enabled transition is enabled for at least δ time units.

The continuous-time semantics of a 1-safe TPN $N = \langle P, T, F, M_0, \alpha, \beta \rangle$ can be conveniently introduced for behaviors over propositions in $\mathcal{P} = \mu \cup \epsilon \cup \tau = \{\mu(p), \epsilon(p) \mid p \in P\} \cup \{\tau(t) \mid t \in T\}$ as follows. Intuitively, at any time t over a be-

havior b , $\mu(p) \in b(t)$ denotes that place p is marked; $\tau(u)$ being triggered at t denotes that transition u fires at t ; and $\epsilon(p)$ being triggered at t denotes that place p undergoes a “zero-time unmarking”, as it will be defined shortly.² Then, b is a *run* of TPN N , and we write $b \models_{\mathbb{R}_{\geq 0}} N$, iff the following conditions hold:

- *Initialization*: $b(0) = \epsilon \cup \tau \cup \bigcup_{p \in M_0} \mu(p)$, and there exists a transition instant $t_{\text{start}} > 0^3$ such that: $b(t) = (t)$ for all $0 \leq t \leq t_{\text{start}}$ and $b^+(t_{\text{start}}) = \tau \cup \bigcup_{p \in M_0} \mu(p)$.
- *Marking*: for all instants $u > t_{\text{start}}$ such that $\mu(p) \notin b^-(u)$ and $\mu(p) \in b^+(u)$ we say that p becomes marked. Correspondingly, there exists a transition $t \in \bullet p$ such that: (i) $\tau(t)$ is triggered at u , (ii) for no other transition $t' \in \bullet p$ (other than t itself) $\tau(t')$ is triggered at u , and (iii) for no transition $t \in p\bullet$ $\tau(t)$ is triggered at u .
- *Unmarking*: for all instants $u > t_{\text{start}}$ such that $\mu(p) \in b^-(u)$ and $\mu(p) \notin b^+(u)$ we say that p becomes unmarked. Correspondingly, there exists a transition $t \in p\bullet$ such that: (i) $\tau(t)$ is triggered at u , (ii) for no other transition $t' \in p\bullet$ (other than t itself) $\tau(t')$ is triggered at u , and (iii) for no transition $t \in \bullet p$ $\tau(t)$ is triggered at u .
- *Enabling*: for all instants $u > t_{\text{start}}$ such that $\tau(t)$ is triggered at u , all places $p \in \bullet t$ must have been marked continuously over $(u - \alpha(t), u)$ without any zero-time unmarkings of the same places occurring.
- *Bound*: for all instants $u > t_{\text{start}}$ such that $\tau(t)$ has not been triggered anywhere over $(u - \beta(t), u)$ and all places $p \in \bullet t$ have been marked continuously, one of the following must occur: (i) all such p 's becomes unmarked at u , (ii) $\tau(t)$ is triggered at u , or (iii) all such p 's are still marked “now on” and some $p \in \bullet t$ undergoes a zero-time unmarking (i.e., $\epsilon(p)$ is triggered at u).
- *Effect*: for all instants $u > t_{\text{start}}$ such that $\tau(t)$ is triggered at u , any place $p \in \bullet t$ becomes unmarked or undergoes a zero-time unmarking, and any place $p \in t\bullet$ becomes marked or undergoes a zero-time unmarking.
- *Zero-time unmarking*: for all instants $u > t_{\text{start}}$ such that $\epsilon(p)$ is triggered at u we say that p undergoes a zero-time unmarking. Correspondingly, there exist transitions $t_a \in \bullet p$ and $t_b \in p\bullet$ such that $\tau(t_a)$ is triggered, $\tau(t_b)$ is triggered, and for no other transition $t' \in \bullet p \cup p\bullet$ (other than t_a, t_b) $\tau(t')$ is triggered.

2.4 Discrete-time approximations of continuous-time specifications

This section provides an overview of the results in [11] that will be used as a basis for the technique of this paper. The technique of [11] is based on two approximation functions for MTL formulas, called under- and over-approximation. The under-approximation function Ω_δ maps continuous-time MTL formulas to discrete-time formulas such that the non-validity of the latter implies the non-validity of the former, over behaviors in \mathcal{B}_χ^δ ; in other words Ω_δ preserves validity from continuous to discrete time. The over-approximation function O_δ maps continuous-time MTL formulas

²The dual “zero-time markings” do not occur over non-Berkeley behaviors as a consequence of zero-time transitions not occurring.

³In the following, we will assume that $t_{\text{start}} \in [0, 2\delta]$ for the discretization parameter $\delta > 0$.

to discrete-time MTL formulas such that the validity of the latter implies the validity of the former, over behaviors in \mathcal{B}_χ^δ . We have the following fundamental verification result, which constitutes the basis of the whole verification framework in the paper.

Proposition 1 (Approximations [11]). *For any MTL formulas ϕ_1, ϕ_2 , and for any $\delta \in \mathcal{D}_{\phi_1, \phi_2}$: (1) if $\text{Alw}(\Omega_\delta(\phi_1)) \Rightarrow \text{Alw}(\text{O}_\delta(\phi_2))$ is \mathbb{N} -valid, then $\text{Alw}(\phi_1) \Rightarrow \text{Alw}(\phi_2)$ is χ^δ -valid; and (2) if $\text{Alw}(\text{O}_\delta(\phi_1)) \Rightarrow \text{Alw}(\Omega_\delta(\phi_2))$ is not \mathbb{N} -valid, then $\text{Alw}(\phi_1) \Rightarrow \text{Alw}(\phi_2)$ is not χ^δ -valid.*

Proposition 1 suggests the following verification approach for MTL. Assume first a system modeled as an (arbitrarily complex) MTL formula ϕ^{sys} ; in order to verify if another MTL formula ϕ^{prop} holds for all run of the system we should check the *validity* of the derived MTL formula $\text{Alw}(\phi^{\text{sys}}) \Rightarrow \text{Alw}(\phi^{\text{prop}})$ which postulates that every run of the system also satisfies the property. Over continuous time, we would build the two discrete-time formulas of Proposition 1 and infer the validity of the continuous-time formula from the results of a discrete-time validity checking. The technique is incomplete as, in particular, when approximation (1) is not valid and approximation (2) is valid nothing can be inferred about the validity of the property in the original system over continuous time.

Consider now another notation \mathcal{N} (e.g., TA or TPN); if we can characterize the continuous-time semantics of any system described with \mathcal{N} by means of a set of MTL formulas, we can reduce the (continuous-time) verification problem for \mathcal{N} to the (continuous-time) verification problem for MTL, and solve the latter as outlined in the previous paragraph.

There are, however, several practical hurdles that make this approach not straightforward to achieve. First, the application of the over- and under- approximations of [11] requires MTL formulas written in a particular form and which do not nest temporal operators. Although in principle every formula can be transformed in the required form (possibly with the addition of a finite number of fresh propositional variables), not any transformation is effective. That is, it turns out that semantically equivalent continuous-time formulas can yield dramatically different — in terms of efficacy and completeness — approximated discrete-time formulas. The axiomatization of operational formalisms (such as TA and TPN) is all the more extremely tricky and requires different sets of axioms, according to whether they will undergo under- or over- approximation. However, all different axiomatizations will be shown to be continuous-time equivalent, hence the intended semantics is captured correctly in all situations. The application in practice of the MTL verification technique will use the “best” set of axioms in every case.

3 Discretizable MTL Axiomatizations of TPN

It is not too hard to devise a general, continuous-time axiomatization of the semantics of a non-trivial subclass of TPN. However, this axiomatization—for reasons that are similar to those discussed in [12] for the TA axiomatization—yields a poor discretized counterpart when the technique of Section 2.4 is applied. Then, this section describes three equivalent (for non-Berkeley behaviors) continuous-time axiomatizations of the semantics of TPN (as introduced in Section 2.3): a generic one (Section 3.1), one that works best for discrete-time under-approximation (Section 3.2), and one that works best for discrete-time over-approximation (Section 3.4). Sections 3.3 and 3.5 produce

respectively the corresponding discrete-time formulas that will be used in the verification problem. Throughout this section, assume a TPN $N = \langle P, T, F, M_0, \alpha, \beta \rangle$ and the set of propositions $\mathcal{P} = \mu \cup \epsilon \cup \tau$ as in the definition of their semantics (Section 2.3). The axiomatization of TPN presented in this paper imposes that, in every marking, a place can contain at most one token. As a consequence, it captures all evolutions of any TPN that is 1-safe; however, it is also capable of describing, for a TPN that is not 1-safe (i.e., which has reachable markings such that at least one place contains more than one token) the sequences of markings in which every place has at most one token. For 1-safe TPN (either by construction or by imposition) any marking M is completely described by the subset of places that are marked in M , which simplifies their formalization. We remark, however, that extending the axiomatization to include generic *bounded* TPN would be routine.

3.1 Generic axiomatization

The continuous-time semantics of a 1-safe TPN $N = \langle P, T, F, M_0, \alpha, \beta \rangle$ can be described through the set of propositions $\mathcal{P} = \mu \cup \epsilon \cup \tau$, where $\mu = \{\mu_p \mid p \in P\}$, $\epsilon = \{\epsilon_p \mid p \in P\}$ and $\tau = \{\tau_u \mid u \in T\}$. Intuitively, at any time t in a behavior b , $\mu_p \in b(t)$ denotes that place p is marked; τ_u being “triggered” (see Section 2) at t denotes that transition u fires at t ; and ϵ_p being triggered at t denotes that place p undergoes a “zero-time unmarking”, that is, p is both unmarked and marked at the same instant (hence does not change the number of contained tokens), as it will be defined shortly.⁴ Then, b is a *run* of TPN N , and we write $b \models_{\mathbb{R}_{\geq 0}} N$, iff the conditions listed below hold.

3.1.1 Places

Marking and unmarking of place $p \in P$ is described by linking transitions of μ_p to transitions of τ_u for transitions u in the pre and postset of p . The trigger operator \mathfrak{I} (matching Δ) is used for τ_u as the actual truth value of τ_u after the transition is irrelevant as long as a transition occurs.

Marking: For all instants t such that μ_p becomes true in t we say that p becomes marked. Correspondingly, there exists a transition $u \in \bullet p$ such that: (i) τ_u is triggered at t , (ii) for no other transition $u' \in \bullet p$ (other than u itself) $\tau_{u'}$ is triggered at t , and (iii) for no transition $u'' \in p \bullet$ ($\tau(u'')$ is triggered at t). This corresponds to the following axioms.

$$p \in M_0 : \Delta(\mu_p) \Rightarrow \left(\bigvee_{u \in \bullet p} \left(\mathfrak{I}(\tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\mathfrak{I}}(\tau_{u'}) \right) \wedge \bigwedge_{u \in p \bullet} \bar{\mathfrak{I}}(\tau_u) \right) \vee \bar{\square}_{(0, \infty)}(\neg \mu_p) \quad (1)$$

$$p \notin M_0 : \Delta(\mu_p) \Rightarrow \bigvee_{u \in \bullet p} \left(\mathfrak{I}(\tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\mathfrak{I}}(\tau_{u'}) \right) \wedge \bigwedge_{u \in p \bullet} \bar{\mathfrak{I}}(\tau_u) \quad (2)$$

⁴The dual “zero-time markings” (in which a place p is both marked and unmarked at the same instant, and hence remains empty) do not occur over non-Berkeley behaviors since, over these behaviors, transitions cannot fire in the same instant in which they are enabled.

Unmarking: For all instants t such that μ_p becomes false in t we say that p becomes unmarked. Correspondingly, there exists a transition $u \in p\bullet$ such that: (i) τ_u is triggered at t , (ii) for no other transition $u' \in p\bullet$ (other than u itself) $\tau_{u'}$ is triggered at t , and (iii) for no transition $u'' \in \bullet p$ $\tau(u'')$ is triggered at t .

$$\Delta(\neg\mu_p) \Rightarrow \bigvee_{u \in p\bullet} \left(\mathfrak{I}(\tau_u) \wedge \bigwedge_{u' \neq u \in p\bullet} \bar{\mathfrak{I}}(\tau_{u'}) \right) \wedge \bigwedge_{u \in \bullet p} \bar{\mathfrak{I}}(\tau_u) \quad (3)$$

3.1.2 Transitions

The lower and upper bounds on the firing of transition u are specified by necessary and sufficient conditions, respectively, on transitions of proposition τ_u . Earliest and latest firing times are introduced through MTL real-time constraints. A non-firing transition u stays enabled as long as μ_p (for p in t 's preset) holds continuously.

Enabling: For all instants t such that τ_u is triggered at t , all places $p \in \bullet u$ must have been marked continuously over $(t - \alpha(u), t)$ without any zero-time unmarkings of the same places occurring.

$$\mathfrak{I}(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \widetilde{\text{O}}(\mu_p \wedge \epsilon_p) \wedge \underset{\vee}{\bar{\text{O}}}_{(0, \alpha(u))}(\mu_p \wedge \epsilon_p) \\ \widetilde{\text{O}}(\mu_p \wedge \neg\epsilon_p) \wedge \underset{\vee}{\bar{\text{O}}}_{(0, \alpha(u))}(\mu_p \wedge \neg\epsilon_p) \end{array} \right) \quad (4)$$

Bound: For all instants t such that τ_u has not been triggered anywhere over $(t - \beta(u), t)$ and all places $p \in \bullet u$ have been marked continuously, one of the following must occur: (i) one of such p 's becomes unmarked at t , (ii) τ_u is triggered at t , or (iii) all such p 's are still marked in $b^+(t)$ and some $p \in \bullet u$ undergoes a zero-time unmarking (i.e., ϵ_p is triggered at t). This is formalized by introducing two axioms for each transition $u \in T$.

$$\bar{\text{O}}_{(0, \beta(u))} \left(\tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} (\neg\mu_p \vee \widetilde{\text{O}}(\neg\mu_p)) \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \bar{\text{O}}_{(0, \beta(u))}(\epsilon_p) \Rightarrow \neg\epsilon_p \vee \widetilde{\text{O}}(\neg\epsilon_p) \\ \bar{\text{O}}_{(0, \beta(u))}(\neg\epsilon_p) \Rightarrow \epsilon_p \vee \widetilde{\text{O}}(\epsilon_p) \end{array} \right) \\ \neg\tau_u \vee \widetilde{\text{O}}(\neg\tau_u) \end{array} \right) \quad (5)$$

$$\bar{\text{O}}_{(0, \beta(u))} \left(\neg\tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} (\neg\mu_p \vee \widetilde{\text{O}}(\neg\mu_p)) \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \bar{\text{O}}_{(0, \beta(u))}(\epsilon_p) \Rightarrow \neg\epsilon_p \vee \widetilde{\text{O}}(\neg\epsilon_p) \\ \bar{\text{O}}_{(0, \beta(u))}(\neg\epsilon_p) \Rightarrow \epsilon_p \vee \widetilde{\text{O}}(\epsilon_p) \end{array} \right) \\ \tau_u \vee \widetilde{\text{O}}(\tau_u) \end{array} \right) \quad (6)$$

Axioms (5–6) impose a so-called “strong time semantics” to the TPN model [8]. This is a departure from the notion of TA formalized in [12], for which the axioms impose what is in fact a weak time semantics [10].

Effect: For all instants t such that τ_u is triggered at t , every place $p \in \bullet u$ either becomes unmarked or undergoes a zero-time unmarking, and every place $p \in u\bullet$ either becomes marked or undergoes a zero-time unmarking.

$$\mathfrak{I}(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} (\Delta(\neg\mu_p) \vee \mathfrak{I}(\epsilon_p)) \wedge \bigwedge_{p \in u\bullet} (\Delta(\mu_p) \vee \mathfrak{I}(\epsilon_p)) \quad (7)$$

3.1.3 Zero-time unmarking

For all instants t such that ϵ_p is triggered at t we say that p undergoes a zero-time unmarking. Correspondingly, there exist transitions $u_a \in \bullet p$ and $u_b \in p\bullet$ such that τ_{u_a} is triggered, τ_{u_b} is triggered, and for no other transition $u' \in \bullet p \cup p\bullet$ (other than u_a, u_b) $\tau_{u'}$ is triggered.

$$\mathfrak{I}(\epsilon_p) \Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p\bullet}} \left(\begin{array}{c} \mathfrak{I}(\tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \bar{\mathfrak{I}}(\tau_{u'}) \\ \wedge \\ \mathfrak{I}(\tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p\bullet} \bar{\mathfrak{I}}(\tau_{u'}) \end{array} \right) \quad (8)$$

3.1.4 Initialization

$b(0) = \epsilon \cup \tau$, and there exists a transition instant $t_{\text{start}} > 0$ such that: $b(t) = b(0)$ for all $0 \leq t < t_{\text{start}}$ and $b^+(t_{\text{start}}) = \epsilon \cup \tau \cup \bigcup_{p \in M_0} \mu_p$ (i.e., the places in the initial marking become marked at t_{start}). This is captured by the following axiom:

$$\text{at } 0: \bigwedge_{p \in P} \neg\mu_p \wedge \diamond_{[0, 2\delta]} \left(\bigwedge_{p \in M_0} \mu_p \right) \wedge \bigcirc \left(\bigwedge_{p \in P} \epsilon_p \wedge \bigwedge_{u \in T} \tau_u \right) \quad (9)$$

Finally, given a TPN N , the MTL formula ψ_N formalizing N is the conjunction of axioms (1–9) instantiated for each place and transition of N .

3.2 Axiomatization for under-approximation

As also discussed in [12], operator Δ yields very weak under-approximations when used to the left-hand side of implications. It turns out that the under-approximation of $\Delta(\phi_1, \phi_2)$ is the discrete-time formula $\bar{\square}_{[0,1]}(\phi_1) \wedge \phi_2$. For a proposition x , $\Delta(x)$ is then the unsatisfiable formula $\bar{\square}_{[0,1]}(\neg x) \wedge x$; correspondingly all implications with such formulas as antecedent are trivially true and do not constrain in any way the discrete-time system.

The approximations can be significantly improved by using the more constraining \blacktriangle in place of Δ . One can check that the under-approximation of $\blacktriangle(x)$ is $\blacktriangle(x)$ itself, which describes a discrete-time transition with $\neg x$ holding at the current instant and x holding at the next instant. Correspondingly, all instances of Δ are changed into instances of \blacktriangle in (2–9) yielding (11–18).

3.2.1 Places

$$p \in M_0 : \blacktriangle(\mu_p) \Rightarrow \left(\bigvee_{u \in \bullet p} \left(\lambda(\tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\lambda}(\tau_{u'}) \right) \wedge \bigwedge_{u \in p \bullet} \bar{\lambda}(\tau_u) \right) \wedge \bar{\square}_{[0, \infty)}(\neg \mu_p) \quad (10)$$

$$p \notin M_0 : \blacktriangle(\mu_p) \Rightarrow \bigvee_{u \in \bullet p} \left(\lambda(\tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\lambda}(\tau_{u'}) \right) \wedge \bigwedge_{u \in p \bullet} \bar{\lambda}(\tau_u) \quad (11)$$

$$\blacktriangle(\neg \mu_p) \Rightarrow \bigvee_{u \in p \bullet} \left(\lambda(\tau_u) \wedge \bigwedge_{u' \neq u \in p \bullet} \bar{\lambda}(\tau_{u'}) \right) \wedge \bigwedge_{u \in \bullet p} \bar{\lambda}(\tau_u) \quad (12)$$

3.2.2 Transitions

$$\lambda(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \mu_p \wedge \epsilon_p \wedge \bar{\square}_{(0, \alpha(u) - \delta)}(\mu_p \wedge \epsilon_p) \\ \vee \\ \mu_p \wedge \neg \epsilon_p \wedge \bar{\square}_{(0, \alpha(u) - \delta)}(\mu_p \wedge \neg \epsilon_p) \end{array} \right) \quad (13)$$

$$\text{Same as (5)} \quad (14)$$

$$\text{Same as (6)} \quad (15)$$

$$\lambda(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\blacktriangle(\neg \mu_p) \vee \lambda(\epsilon_p) \right) \wedge \bigwedge_{p \in u \bullet} \left(\blacktriangle(\mu_p) \vee \lambda(\epsilon_p) \right) \quad (16)$$

3.2.3 Zero-time unmarking

$$\lambda(\epsilon_p) \Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p \bullet}} \left(\begin{array}{c} \lambda(\tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \bar{\lambda}(\tau_{u'}) \\ \wedge \\ \lambda(\tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p \bullet} \bar{\lambda}(\tau_{u'}) \end{array} \right) \quad (17)$$

3.2.4 Initialization

$$\bar{\square}_{[\delta, \infty)}(\perp) \Rightarrow \bigwedge_{p \in P} \neg \mu_p \wedge \diamond_{[0, 2\delta]} \left(\bigwedge_{p \in M_0} \mu_p \right) \wedge \circ \left(\bigwedge_{p \in P} \epsilon_p \wedge \bigwedge_{u \in T} \tau_u \right) \quad (18)$$

It can be shown that (1–9) are equivalent to (10–18) over behaviors that are non-Berkeley for δ . For instance, consider (2) and (11). In order to show that (2) implies (11), let (2) and $\blacktriangle(\mu_p)$ hold at the current time instant z . $\blacktriangle(\mu_p)$ implies that there exists a $z' \in [z, z + \delta]$ where μ_p shifts from false to true. (2) evaluated at z' entails (among other things) that $\lambda(\tau_u)$ holds at z' for some t ; that is, τ_u is triggered at z' . Without loss of generality, assume that τ_u is false before z' and is true after it. The non-Berkeleyness assumption allows us to strengthen this fact, so that τ_u is false at z as well and is true until $z + \delta$, because $z' \in [z, z + \delta]$. Hence $\lambda(\tau_u)$ holds at z . The rest of the implication is proved similarly. The proof of the converse implication that (11) implies (2) also relies on the non-Berkeleyness assumption, which guarantees that there is exactly one transition of μ_p over $[z, z + \delta]$ as a consequence of $\blacktriangle(\mu_p)$ holding at z . We omit the details of the proof, which are however along the same lines.

3.3 Under-approximation

The under-approximations of (10–18) are reported as formulas (19–27). Notice the lower- and upper-bound relaxations in (22–24), in accordance with the notion of under-approximation.

3.3.1 Places

Syntactically the same as in (10) (19)

Syntactically the same as in (11) (20)

Syntactically the same as in (12) (21)

3.3.2 Transitions

$$\lambda(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \mu_p \wedge \epsilon_p \wedge \overline{\square}_{[1, \alpha(u)/\delta - 2]}(\mu_p \wedge \epsilon_p) \\ \mu_p \wedge \neg \epsilon_p \wedge \overline{\square}_{[1, \alpha(u)/\delta - 2]}(\mu_p \wedge \neg \epsilon_p) \end{array} \right) \quad (22)$$

$$\overline{\square}_{[0, \beta(u)/\delta]} \left(\tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} \diamond_{=1}(\neg \mu_p) \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \overline{\square}_{[0, \beta(u)/\delta]}(\epsilon_p) \Rightarrow \diamond_{=1}(\neg \epsilon_p) \\ \overline{\square}_{[0, \beta(u)/\delta]}(\neg \epsilon_p) \Rightarrow \diamond_{=1}(\epsilon_p) \end{array} \right) \\ \diamond_{=1}(\neg \tau_u) \end{array} \right) \quad (23)$$

$$\overline{\square}_{[0, \beta(u)/\delta]} \left(\neg \tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} \diamond_{=1}(\neg \mu_p) \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \overline{\square}_{[0, \beta(u)/\delta]}(\epsilon_p) \Rightarrow \diamond_{=1}(\neg \epsilon_p) \\ \overline{\square}_{[0, \beta(u)/\delta]}(\neg \epsilon_p) \Rightarrow \diamond_{=1}(\epsilon_p) \end{array} \right) \\ \diamond_{=1}(\tau_u) \end{array} \right) \quad (24)$$

Syntactically the same as in (16) (25)

The straightforward under-approximation of (14) and (15) yields formulas which have been re-arranged to eliminate redundant terms. In fact, the time bound $(0, \beta(u))$ in the antecedent becomes $[0, \beta(u)/\delta]$ when under-approximated. Hence, formulas such as $\overline{\square}_{(0, \beta(u))}(\gamma) \Rightarrow \neg \gamma \vee \bigcirc(\neg \gamma)$ are under-approximated as $\overline{\square}_{[0, \beta(u)/\delta]}(\gamma) \Rightarrow \neg \gamma \vee \diamond_{[0, 1]}(\neg \gamma)$. However, $\neg \gamma$ never holds at the current instant because it would contradict the antecedent. Correspondingly, such formulas can be simplified to $\overline{\square}_{[0, \beta(u)/\delta]}(\gamma) \Rightarrow \diamond_{=1}(\neg \gamma)$.

3.3.3 Zero-time unmarking

Syntactically the same as in (17) (26)

3.3.4 Initialization

$$\text{at } 0: \bigwedge_{p \in P} \neg \mu_p \wedge \diamond_{[1, 2]} \left(\bigwedge_{p \in M_0} \mu_p \right) \wedge \bigwedge_{p \in P} \epsilon_p \wedge \bigwedge_{u \in T} \tau_u \quad (27)$$

3.4 Axiomatization for over-approximation

Continuous-time operator Δ becomes⁵ discrete-time operator \blacktriangle under over-approximation when it occurs to the left-hand side of implications, hence is suitable to describe antecedents of transitions that will be over-approximated. However, the over-approximation of the same operator takes a different form in the right-hand side of implications. In such cases, the over-approximation of formulas such as $\Delta(x)$ is $\overline{\square}_{[0,1]}(\neg x) \wedge \square_{[0,1]}(x)$ which is clearly unsatisfiable. Correspondingly, the whole over-approximation formulas would be unsatisfiable only for false antecedents, i.e., when no transition ever occurs.

After careful experimentation, we found that a workaround to this problem should exploit a weakening of the Δ operators that occur in consequent formulas. Let us illustrate the idea as simply as possible for two propositions x, y and the formula $\Delta(x) \Rightarrow \Delta(y)$: every transition of x occurs concurrently with a transition of y . The formula is relaxed into the weaker $\Delta(x) \Rightarrow \widetilde{\bigcirc}(\neg y) \wedge \square_{=\delta}(x \Rightarrow y)$: every transition of x also triggers a transition of y sometime in the future, as long as x still holds δ time units in the future. The new formula is essentially equivalent to the original one for non-Berkeley behaviors for the following reasons. First, x must still hold δ time units in the future, because its behavior is non-Berkeley for δ ; hence y holds as well there and must transition somewhere over the interval $(0, \delta)$ from the current instant. In addition, the transition of y cannot occur asynchronously to the transition of x ; otherwise two distinct transitions would occur within δ time units, against the non-Berkeleyness assumption. In all, the two formulations are equivalent over non-Berkeley continuous time. Correspondingly, the \mathbb{W} operator is introduced and used in the right-hand side of implications in the following continuous-time formulas (28–36).

3.4.1 Places

$$p \in M_0 : \Delta(\mu_p) \Rightarrow \left(\begin{array}{c} \bigvee_{u \in \bullet p} \left(\mathbb{W}(\mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \overline{\mathbb{W}}(\mu_p \rightsquigarrow \tau_{u'}) \right) \\ \bigwedge_{u \in p \bullet} \overline{\mathbb{W}}(\mu_p \rightsquigarrow \tau_u) \\ \overline{\square}_{[\delta, \infty)}(\neg \mu_p) \end{array} \right) \quad (28)$$

$$p \notin M_0 : \Delta(\mu_p) \Rightarrow \left(\begin{array}{c} \bigvee_{u \in \bullet p} \left(\mathbb{W}(\mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \overline{\mathbb{W}}(\mu_p \rightsquigarrow \tau_{u'}) \right) \\ \bigwedge_{u \in p \bullet} \overline{\mathbb{W}}(\mu_p \rightsquigarrow \tau_u) \end{array} \right) \quad (29)$$

$$\Delta(\neg \mu_p) \Rightarrow \bigvee_{u \in p \bullet} \left(\mathbb{W}(\neg \mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in p \bullet} \overline{\mathbb{W}}(\neg \mu_p \rightsquigarrow \tau_{u'}) \right) \wedge \bigwedge_{u \in \bullet p} \overline{\mathbb{W}}(\neg \mu_p \rightsquigarrow \tau_u) \quad (30)$$

⁵After some semantic-preserving simplifications.

3.4.2 Transitions

$$\mathbb{I}(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \widetilde{\mathbb{O}}(\mu_p \wedge \epsilon_p) \wedge \widetilde{\mathbb{I}}_{[\delta, \alpha(u))}(\mu_p \wedge \epsilon_p) \\ \widetilde{\mathbb{O}}(\mu_p \wedge \neg \epsilon_p) \wedge \widetilde{\mathbb{I}}_{[\delta, \alpha(u))}(\mu_p \wedge \neg \epsilon_p) \end{array} \right) \quad (31)$$

$$\text{Same as (5)} \quad (32)$$

$$\text{Same as (6)} \quad (33)$$

$$\Delta(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \left(\begin{array}{c} \widetilde{\mathbb{O}}(\mu_p) \\ \square_{=\delta}(\tau_u \Rightarrow \neg \mu_p) \\ \mathbb{I}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \\ \left(\begin{array}{c} \widetilde{\mathbb{O}}(\neg \mu_p) \\ \square_{=\delta}(\tau_u \Rightarrow \mu_p) \\ \mathbb{I}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \end{array} \right) \wedge \bigwedge_{p \in u \bullet} \left(\begin{array}{c} \left(\begin{array}{c} \widetilde{\mathbb{O}}(\neg \mu_p) \\ \square_{=\delta}(\tau_u \Rightarrow \mu_p) \\ \mathbb{I}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \\ \left(\begin{array}{c} \widetilde{\mathbb{O}}(\mu_p) \\ \square_{=\delta}(\neg \tau_u \Rightarrow \neg \mu_p) \\ \mathbb{I}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \end{array} \right) \quad (34)$$

$$\Delta(\neg \tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \left(\begin{array}{c} \widetilde{\mathbb{O}}(\mu_p) \\ \square_{=\delta}(\neg \tau_u \Rightarrow \neg \mu_p) \\ \mathbb{I}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \\ \left(\begin{array}{c} \widetilde{\mathbb{O}}(\neg \mu_p) \\ \square_{=\delta}(\neg \tau_u \Rightarrow \mu_p) \\ \mathbb{I}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \end{array} \right) \wedge \bigwedge_{p \in u \bullet} \left(\begin{array}{c} \left(\begin{array}{c} \widetilde{\mathbb{O}}(\neg \mu_p) \\ \square_{=\delta}(\neg \tau_u \Rightarrow \mu_p) \\ \mathbb{I}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \\ \left(\begin{array}{c} \widetilde{\mathbb{O}}(\mu_p) \\ \square_{=\delta}(\tau_u \Rightarrow \neg \mu_p) \\ \mathbb{I}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \end{array} \right) \quad (34)$$

3.4.3 Zero-time unmarking

$$\begin{aligned} \Delta(\epsilon_p) &\Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p \bullet}} \left(\begin{array}{c} \mathbb{I}(\epsilon_p \rightsquigarrow \tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \widetilde{\mathbb{I}}(\epsilon_p \rightsquigarrow \tau_{u'}) \\ \mathbb{I}(\epsilon_p \rightsquigarrow \tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p \bullet} \widetilde{\mathbb{I}}(\epsilon_p \rightsquigarrow \tau_{u'}) \end{array} \right) \\ \Delta(\neg \epsilon_p) &\Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p \bullet}} \left(\begin{array}{c} \mathbb{I}(\neg \epsilon_p \rightsquigarrow \tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \widetilde{\mathbb{I}}(\neg \epsilon_p \rightsquigarrow \tau_{u'}) \\ \mathbb{I}(\neg \epsilon_p \rightsquigarrow \tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p \bullet} \widetilde{\mathbb{I}}(\neg \epsilon_p \rightsquigarrow \tau_{u'}) \end{array} \right) \end{aligned} \quad (35)$$

3.4.4 Initialization

$$\widetilde{\mathbb{I}}_{(0, \infty)}(\perp) \Rightarrow \bigwedge_{p \in P} \neg \mu_p \wedge \diamond_{[0, 2\delta]} \left(\bigwedge_{p \in M_0} \mu_p \right) \wedge \bigcirc \left(\bigwedge_{p \in P} \epsilon_p \wedge \bigwedge_{u \in T} \tau_u \right) \quad (36)$$

The observations that have been introduced at the beginning of this section can be leveraged to provide a rigorous proof that (28–36) are equivalent to the original (1–9) over non-Berkeley continuous time. We omit the details for brevity.

3.5 Over-approximation

The over-approximations of (28–36) are reported as formulas (37–45). Notice the lower- and upper-bound relaxations in (40–42), in accordance with the notion of over-approximation.

3.5.1 Places

$$p \in M_0 : \blacktriangle(\mu_p) \Rightarrow \left(\begin{array}{c} \bigvee_{u \in \bullet p} \left(\mathbb{M}(\mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\mathbb{M}}(\mu_p \rightsquigarrow \tau_{u'}) \right) \\ \bigwedge_{u \in p \bullet} \bar{\mathbb{M}}(\mu_p \rightsquigarrow \tau_u) \\ \bar{\square}_{[\delta, \infty)}(\neg \mu_p) \end{array} \right) \quad (37)$$

$$p \notin M_0 : \blacktriangle(\mu_p) \Rightarrow \left(\begin{array}{c} \bigvee_{u \in \bullet p} \left(\mathbb{M}(\mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in \bullet p} \bar{\mathbb{M}}(\mu_p \rightsquigarrow \tau_{u'}) \right) \\ \bigwedge_{u \in p \bullet} \bar{\mathbb{M}}(\mu_p \rightsquigarrow \tau_u) \end{array} \right) \quad (38)$$

$$\blacktriangle(\neg \mu_p) \Rightarrow \bigvee_{u \in p \bullet} \left(\mathbb{M}(\neg \mu_p \rightsquigarrow \tau_u) \wedge \bigwedge_{u' \neq u \in p \bullet} \bar{\mathbb{M}}(\neg \mu_p \rightsquigarrow \tau_{u'}) \right) \wedge \bigwedge_{u \in \bullet p} \bar{\mathbb{M}}(\neg \mu_p \rightsquigarrow \tau_u) \quad (39)$$

3.5.2 Transitions

$$\mathfrak{R}(\tau_u) \Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \bar{\square}_{[0, \alpha(u)/\delta + 1]}(\mu_p \wedge \epsilon_p) \\ \bigvee \\ \bar{\square}_{[0, \alpha(u)/\delta + 1]}(\mu_p \wedge \neg \epsilon_p) \end{array} \right) \quad (40)$$

$$\bar{\square}_{[1, \beta(u)/\delta - 1]} \left(\tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} \left(\neg \mu_p \vee \square_{[0, 1]}(\neg \mu_p) \right) \\ \bigvee \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \bar{\square}_{[1, \beta(u)/\delta - 1]}(\epsilon_p) \Rightarrow \neg \epsilon_p \vee \square_{[0, 1]}(\neg \epsilon_p) \\ \bigwedge \\ \bar{\square}_{[0, \beta(u)/\delta - 1]}(\neg \epsilon_p) \Rightarrow \epsilon_p \vee \diamond_{[0, 1]}(\epsilon_p) \end{array} \right) \\ \bigvee \\ \neg \tau_u \end{array} \right) \quad (41)$$

$$\bar{\square}_{[1, \beta(u)/\delta - 1]} \left(\neg \tau_u \wedge \bigwedge_{p \in \bullet u} \mu_p \right) \Rightarrow \left(\begin{array}{c} \bigvee_{p \in \bullet u} \square_{[0, 1]}(\neg \mu_p) \\ \bigvee \\ \bigvee_{p \in \bullet u} \left(\begin{array}{c} \bar{\square}_{[1, \beta(u)/\delta - 1]}(\epsilon_p) \Rightarrow \neg \epsilon_p \vee \square_{[0, 1]}(\neg \epsilon_p) \\ \bigwedge \\ \bar{\square}_{[0, \beta(u)/\delta - 1]}(\neg \epsilon_p) \Rightarrow \epsilon_p \vee \diamond_{[0, 1]}(\epsilon_p) \end{array} \right) \\ \bigvee \\ \tau_u \end{array} \right) \quad (42)$$

$$\begin{aligned} \blacktriangle(\tau_u) &\Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \bar{\square}_{[0, 1]}(\mu_p) \\ \bigwedge \\ \square_{[0, 2]}(\tau_u \Rightarrow \neg \mu_p) \\ \bigvee \\ \mathbb{M}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \wedge \bigwedge_{p \in u \bullet} \left(\begin{array}{c} \bar{\square}_{[0, 1]}(\neg \mu_p) \\ \bigwedge \\ \square_{[0, 2]}(\tau_u \Rightarrow \mu_p) \\ \bigvee \\ \mathbb{M}(\tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \\ \blacktriangle(\neg \tau_u) &\Rightarrow \bigwedge_{p \in \bullet u} \left(\begin{array}{c} \bar{\square}_{[0, 1]}(\mu_p) \\ \bigwedge \\ \square_{[0, 2]}(\neg \tau_u \Rightarrow \neg \mu_p) \\ \bigvee \\ \mathbb{M}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \wedge \bigwedge_{p \in u \bullet} \left(\begin{array}{c} \bar{\square}_{[0, 1]}(\neg \mu_p) \\ \bigwedge \\ \square_{[0, 2]}(\neg \tau_u \Rightarrow \mu_p) \\ \bigvee \\ \mathbb{M}(\neg \tau_u \rightsquigarrow \epsilon_p) \end{array} \right) \end{aligned} \quad (43)$$

Similarly as with under-approximation, formulas have been conveniently simplified: the term $\bar{\square}_{[0, 1]}(\mu_p \wedge \epsilon_p)$ in the consequent of (31) is over-approximated to $\bar{\square}_{[0, 1]}(\mu_p \wedge \epsilon_p)$, which is subsumed by the other term $\bar{\square}_{[0, \alpha(u)/\delta + 1]}(\mu_p \wedge \epsilon_p)$ in the over-approximation. (In fact, $\alpha(u)/\delta + 1 \geq 2$ is the case). Subformulas $\neg \tau_u \vee \square_{[0, 1]}(\neg \tau_u)$ and $\tau_u \vee \square_{[0, 1]}(\tau_u)$ in the over-approximations (42) and (42), respectively, can also be simplified. In fact, (40) enforces marking and no zero-time unmarking for at least 3 time units whenever τ_u is triggered; hence μ_p cannot be triggered over $[0, 1]$ so that the terms $\square_{[0, 1]}(\neg \tau_u)$ and $\square_{[0, 1]}(\tau_u)$ are redundant.

3.5.3 Zero-time unmarking

$$\begin{aligned}
\blacktriangle(\epsilon_p) &\Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p \bullet}} \left(\begin{array}{c} \mathbb{M}(\epsilon_p \rightsquigarrow \tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \bar{\mathbb{M}}(\epsilon_p \rightsquigarrow \tau_{u'}) \\ \wedge \\ \mathbb{M}(\epsilon_p \rightsquigarrow \tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p \bullet} \bar{\mathbb{M}}(\epsilon_p \rightsquigarrow \tau_{u'}) \end{array} \right) \\
\blacktriangle(\neg\epsilon_p) &\Rightarrow \bigvee_{\substack{u_a \in \bullet p \\ u_b \in p \bullet}} \left(\begin{array}{c} \mathbb{M}(\neg\epsilon_p \rightsquigarrow \tau_{u_a}) \wedge \bigwedge_{u' \neq u_a \in \bullet p} \bar{\mathbb{M}}(\neg\epsilon_p \rightsquigarrow \tau_{u'}) \\ \wedge \\ \mathbb{M}(\neg\epsilon_p \rightsquigarrow \tau_{u_b}) \wedge \bigwedge_{u' \neq u_b \in p \bullet} \bar{\mathbb{M}}(\neg\epsilon_p \rightsquigarrow \tau_{u'}) \end{array} \right)
\end{aligned} \tag{44}$$

3.5.4 Initialization

$$\text{at } 0: \bigwedge_{p \in P} \neg\mu_p \wedge \diamond_{=1} \left(\bigwedge_{p \in M_0} \mu_p \right) \wedge \square_{[0,1]} \left(\bigwedge_{p \in P} \epsilon_p \wedge \bigwedge_{u \in T} \tau_u \right) \tag{45}$$

3.6 Quality of discrete-time approximations

Proposition 1 guarantees that under-approximations preserve validity and over-approximations preserve counterexamples. It does not say anything about the *quality* (or completeness) of such approximations; in particular an under-approximation can preserve validity trivially by being contradictory (i.e., inconsistent), and an over-approximation can preserve counterexamples trivially by being identically valid.

In order to make sure this is not the case, let us introduce a set of constraints that guarantees no degenerate behaviors are modeled in the approximations. Consider formulas involving metric intervals, namely (22–24) for the under-approximations and (40–42) for the over-approximation. We should check that, for every transition u with dense-time firing interval $[\alpha(u), \beta(u)]$:

- *non-emptiness*. Metric intervals are non-empty; that is $\alpha(u) \geq 3\delta$ from the under-approximation and $\alpha(u) \geq -\delta$, $\beta(u) \geq 2\delta$ from the over-approximation.
- *consistency*. The the minimum enabling interval (defined in (22) and (40) for under- and over-approximation respectively) is smaller than the maximum enabling interval (defined in (23–24) and (41–42) for under- and over-approximation respectively). Correspondingly, we have the constraints $\beta(u) \geq \alpha(u) - 2\delta$ from the under-approximation and $\beta(u) \geq \alpha(u) + 2\delta$ from the over-approximation.

The constraints can be summarized as $\alpha(u) \geq 3\delta$ and $\beta(u) \geq \alpha(u) + 2\delta$. In our examples, we will consider only non-degenerate TPN satisfying these constraints.

4 Multi-Paradigm Modeling and Verification at Work

The multi-paradigm modeling technique presented in this paper is supported by the *Zot* bounded satisfiability checker [16, 17]. More precisely, we exploited the flexibility provided by the SAT-based approach pursued by *Zot*, and implemented several separate plugins to deal with the various allowed formalisms. In particular, the tool now includes plugins capable of dealing with dense-time MTL formulas [11], with timed automata [12], and with timed Petri nets (using the formalization presented in Section 3). In addition, *Zot* is natively capable of accepting discrete-time MTL formulas as

input language. The plugins provide primitives through which the user can define the system to be analyzed as a mixture of timed automata, dense- and discrete-time MTL formulas, and timed Petri nets. The properties to be verified for the system can also be described as a combination of fragments written using the aforementioned formal languages, though they are usually formalized through MTL formulas (either using dense or discrete time).

The tool then automatically builds, for the dense-time fragments of the system and of the property to be analyzed, the two discrete-time approximation formulas of Proposition 1. These formulas, in possibly conjunction with MTL formulas natively written using a discrete notion of time, are checked for validity over time \mathbb{N} ; the results of the validity check allows one to infer the validity of the integrated model, according to Proposition 1.

The multi-paradigm verification process in *Zot* consists of three sequential phases. First, the discrete-time MTL formulas of Proposition 1 are built and are translated into a propositional satisfiability (SAT) problem. Second, the SAT instance (possibly including MTL formulas directly written using a discrete notion of time) is put into conjunctive normal form (CNF), a standard input format for SAT solvers. Third, the CNF formula is fed to a SAT solving engine (such as MiniSat, zChaff, or MiraXT).

4.1 An Example of Multi-paradigm Modeling and Verification

We demonstrate how the modeling and verification technique presented in this paper works in practice through an example consisting of a fragment of a realistic monitoring system, which could be part of a larger supervision and control system.

The monitoring subsystem is composed of three identical sensors, a middle component that is in charge of acquiring and pre-processing the data from the sensors, and a data management component that further elaborates the data (e.g., to select appropriate control actions). For reasons of dependability (by redundancy), the three sensors measure the same quantity (whose nature is of no relevance in this example). Each one of them senses independently the measured quantity at a certain rate which is in general aperiodic; however, while the acquisition rate can vary, the distance between consecutive acquisitions must always be no less than $T/2$ and no more than T time units. Each sensor keeps track of only the last measurement, hence every new sensed value replaces the one stored by the sensor.

The data acquisition component retrieves data from the three sensors in a “pull” fashion. More precisely, when all three sensors have a fresh measurement available, with a delay of at least $T/10$ units, but of no more than $T/5$ time units, the data acquisition component collects the three values from the sensors (which then become stale, as they have been acquired). After having retrieved the three measurements, the component processes them (e.g., it computes a derived measurement as the average of the sensed values); the process takes between $T/5$ and $T/2$ time units.

After having computed the derived measurement, the data acquisition component sends it to the data manager, this time using a “push” policy which requires an acknowledgement of the data reception by the latter. The data acquisition component tries to send data to the data manager at most twice. If both attempts at data transmission fail (for example because a timeout for the reception acknowledgement by the data manager expires, or because the latter signals a reception error), the data transmission terminates with an error.

First, we model the mechanism through which the three sensors collect data from the field and the data acquisition component retrieves them for the pre-processing

sen, with a certain degree of arbitrariness, because it was deemed capable of representing the timing constraints on the protocol in a more natural way, especially for what concerns the constraint on the overall duration of the process.

Finally, MTL formulas are added to “bridge the gap” between the fragments shown in Figures 1 and 2. This is achieved by the two following formulas, which define, respectively, that the transmission procedure can begin only if a pre-processed measurement value has been produced by the data acquisition component in the last T time units (46) and if the system is not in the middle of a data transmission (i.e., it is idle), and a new datum is being processed, a transmission will start within $T/2$ time units, due to the upper bound of *process.d* transition (47).

$$\text{try} \Rightarrow \overleftarrow{\diamond}_{(0, T/2]}(\text{data_retrieved}) \quad (46)$$

$$\text{data_retrieved} \wedge \text{idle} \Rightarrow \diamond_{(0, T/2]}(\text{try}) \quad (47)$$

Notice that the automata of Figures 1 and 2 are defined, as per the formalizations of [12] and of Section 3, over a continuous notion of time. This choice for the time domain of these two system fragments is justified by the fact that they deal with parts of the system interacting with physical elements (measured quantities, transmission channel), for which a continuous time seems better suited.

Formulas (46) and (47), instead, describe a software synchronization mechanism within the application. As a consequence, discrete time is more suitable to describe this part of the system, hence formulas (46) and (47) are to be interpreted accordingly.

Finally, the model of the system to be verified is built by conjoining the discrete-time approximations for the fragments of Figures 1-2 and the discrete-time MTL formulas (46)-(47). More precisely, if $\psi_N^{\Omega_\delta}$ and $\psi_N^{\text{O}_\delta}$ are the continuous-time MTL formulas capturing the semantics of the net of Figure 1 (see Section 3), $\psi_A^{\Omega_\delta}$, $\psi_A^{\text{O}_\delta}$ are the continuous-time MTL formulas for the automaton of Figure 2, ψ_L is the discrete-time formula $\psi_L = (46) \wedge (47)$, and ϕ^{prop} is the continuous-time property to be checked for the system, then we have:

$$\begin{aligned} \phi^+ &= \text{Alw} \left(\Omega_\delta \left(\psi_N^{\Omega_\delta} \right) \wedge \Omega_\delta \left(\psi_A^{\Omega_\delta} \right) \wedge \psi_L \right) \Rightarrow \text{Alw}(\text{O}_\delta(\phi^{\text{prop}})) \\ \phi^- &= \text{Alw} \left(\text{O}_\delta \left(\psi_N^{\text{O}_\delta} \right) \wedge \text{O}_\delta \left(\psi_A^{\text{O}_\delta} \right) \wedge \psi_L \right) \Rightarrow \text{Alw}(\Omega_\delta(\phi^{\text{prop}})) \end{aligned}$$

Note that formula ψ_L , which is to be interpreted over discrete time, must not be approximated. Then, if ϕ^+ is N-valid, we can draw some interesting conclusions.

First, if one implements a continuous-time system that does not vary faster than the sampling time δ (i.e., whose behaviors are in \mathcal{B}_χ^δ), which satisfies ψ_N , ψ_A , and a continuous-time MTL formula ψ' such that $\Omega_\delta(\psi'_L) = \psi_L$, then property ϕ^{prop} holds for this system.

It can be shown that, for any continuous-time MTL formula ϕ , the set of behaviors satisfying $\text{O}_\delta(\phi)$ is a subset of those satisfying $\Omega_\delta(\phi)$ (i.e., $\{b \mid b \models_{\mathbb{N}} \text{O}_\delta(\phi)\} \subseteq \{b \mid b \models_{\mathbb{R}} \Omega_\delta(\phi)\}$). In addition, given a discrete-time behavior b that satisfies $\text{O}_\delta(\phi)$, from [11, Lemma 3] we have that any continuous-time non-Berkeley behavior b' for which b is a sampling satisfies ϕ . Then, any way one reconstructs a continuous-time non-Berkeley behavior b' from a discrete-time one that satisfies $\text{O}_\delta(\phi)$, b' satisfies ϕ . This leads us to conclude that, if one builds a discrete-time system (e.g., a piece of software) which implements — that is, satisfies — $\text{O}_\delta(\psi_N^{\text{O}_\delta})$, $\text{O}_\delta(\psi_A^{\text{O}_\delta})$, ψ_L , this satisfies discrete-time property $\text{O}_\delta(\phi^{\text{prop}})$; in addition, any way one uses a discrete-

time behavior of this system to reconstruct a continuous-time, non-Berkeley behavior, the latter satisfies ψ_N , ψ_A , and ϕ^{prop} .

Finally, if ϕ^- is not \mathbb{N} -valid, a discrete-time system implementing $O_\delta(\psi_N^{\text{O}_\delta})$, $O_\delta(\psi_A^{\text{O}_\delta})$, ψ_L violates property $\Omega_\delta(\phi^{\text{prop}})$.

Verification. We used the system model presented above to check a number of properties to validate the effectiveness of our approach. Table 4.1 shows the results, and duration of the tests. More precisely, for each test the table reports: the checked property; the values of the timing parameters in the model (i.e., T_1, T_2, T_3, T); the temporal bound k of the time domain (as $\mathbb{Z}ot$ is a bounded satisfiability checker, it considers all the behaviors with period $\leq k$); the total amount of time to perform each phase of the verification, namely formula building (including transformation into conjunctive normal form), and propositional satisfiability checking; the results of the tests; the size (in millions of clauses) of the formula fed to the SAT-solver.⁷ Tests were performed instantiating the parameters with different values to get an idea of how the performance of the verification algorithm is affected, both in terms of time to complete the verification and of whether the verification attempt is conclusive. In addition, the timed interaction between the data acquisition and monitoring subsystems is quite subtle and the properties under verification hold in every run of the system only for certain combinations of parameter values. Automated verification allowed us to investigate this fact in some detail.

First, we checked some properties concerning the liveness of the data collection by a sensor X (with $X \in \{1, 2, 3\}$). More precisely, we analyzed whether property (48) holds for the model.⁸

$$\begin{aligned}
& \text{replaceX} \wedge \text{new_dX} \Rightarrow \\
& \diamond_{(0, T+\delta]} (\text{replaceX} \wedge \neg \text{new_dX} \vee \neg \text{replaceX} \wedge \text{new_dX}) \\
& \quad \wedge \\
& \quad \text{replaceX} \wedge \neg \text{new_dX} \Rightarrow \\
& \diamond_{(0, T+\delta]} (\text{replaceX} \wedge \text{new_dX} \vee \neg \text{replaceX} \wedge \neg \text{new_dX}) \\
& \quad \wedge \\
& \quad \neg \text{replaceX} \wedge \text{new_dX} \Rightarrow \\
& \diamond_{(0, T+\delta]} (\neg \text{replaceX} \wedge \neg \text{new_dX} \vee \text{replaceX} \wedge \text{new_dX}) \\
& \quad \wedge \\
& \quad \neg \text{replaceX} \wedge \neg \text{new_dX} \Rightarrow \\
& \diamond_{(0, T+\delta]} (\neg \text{replaceX} \wedge \text{new_dX} \vee \text{replaceX} \wedge \neg \text{new_dX})
\end{aligned} \tag{48}$$

Formula (48) states that triggering events of `replaceX` and `new_dX` transitions must occur within $T + \delta$ (with δ the sampling period) time instants in the future, i.e., that either `replaceX` or `new_dX` must change value within the next $T + \delta$ time instants. The property does not hold in general, since a firing of transition `retrieve_d` would reset the time counters for transitions `replaceX` and `new_dX`. This fact can be pointed out by checking ϕ^- , with $\phi^{\text{prop}} = (48)$, which is unsatisfiable, as shown in Table 4.1.

⁷The verification tool and the complete model used for verification can be found at <http://home.dei.polimi.it/pradella>. Tests have been performed on a PC equipped with two Intel Xeon E5335 Quad-Core Processor 2GHz, 16 Gb of RAM, and GNU/Linux (kernel 2.6.29), using a single core for each test. $\mathbb{Z}ot$ used the SAT-solver MiniSat 2.

⁸Recall that all properties to be proved are implicitly closed with the Alw operator.

If the additional hypothesis that transition `retrieve_d` does not fire along $(0, T + \delta]$, (48) can however be shown to hold. More precisely, if (48) is rewritten, as shown in formula (49), by adding to the antecedents the condition that predicate `retrieve_d` does not change in $(0, T + \delta]$ (i.e., transition `retrieve_d` does not fire in that interval), then the new ϕ^+ is \mathbb{N} -valid (as Table 4.1 shows), hence (49) holds for the system.

$$\begin{aligned}
& \Box_{(0, T + \delta]}(\text{retrieve_d}) \wedge \text{replaceX} \wedge \text{new_dX} \Rightarrow \\
& \Diamond_{(0, T + \delta]}(\text{replaceX} \wedge \neg \text{new_dX} \vee \neg \text{replaceX} \wedge \text{new_dX}) \\
& \quad \wedge \dots \wedge \\
& \Box_{(0, T + \delta]}(\text{retrieve_d}) \wedge \neg \text{replaceX} \wedge \neg \text{new_dX} \Rightarrow \\
& \Diamond_{(0, T + \delta]}(\neg \text{replaceX} \wedge \text{new_dX} \vee \neg \text{replaceX} \wedge \neg \text{new_dX}) \\
& \quad \vee \\
& \Box_{(0, T + \delta]}(\neg \text{retrieve_d}) \wedge \text{replaceX} \wedge \neg \text{new_dX} \Rightarrow \\
& \Diamond_{(0, T + \delta]}(\text{replaceX} \wedge \text{new_dX} \vee \neg \text{replaceX} \wedge \neg \text{new_dX}) \\
& \quad \wedge \dots \wedge \\
& \Box_{(0, T + \delta]}(\neg \text{retrieve_d}) \wedge \neg \text{replaceX} \wedge \neg \text{new_dX} \Rightarrow \\
& \Diamond_{(0, T + \delta]}(\neg \text{replaceX} \wedge \text{new_dX} \vee \neg \text{replaceX} \wedge \neg \text{new_dX})
\end{aligned} \tag{49}$$

Another liveness property is formalized by formula (50), which states that a datum is retrieved (i.e., place `data_retrieved` is marked) at least every $\frac{3T}{2}$ time units.

$$\Diamond_{(0, \frac{3T}{2}]}(\text{data_retrieved}) \tag{50}$$

Property (50) cannot be established with our verification technique as it falls in the incompleteness region (i.e., ϕ^+ is not valid and ϕ^- is valid, as Table 4.1 shows); from the automated check we cannot draw a definitive conclusion on the validity of the property for the system. If, however, the temporal bound of formula (50) is slightly relaxed as in formula (51), not only the verification is conclusive, but it shows that the property in fact holds for the system.

$$\Diamond_{(0, 2T]}(\text{data_retrieved}) \tag{51}$$

Verification also shows that the original formula (50) holds if the bound on transitions `replaceX` of the TPN is changed to $[\frac{4T}{5}, T]$ (property (50') in Table 4.1).

Formula (52) expresses the maximum delay between sensor collect and data send. More precisely, if each sensor has provided a measurement and transition `retrieve_d` fires, then the timed automaton will enter state `try` within T instants. The validity of this formula would allow us to check that the two parts of the system modeled by the TPN and by the TA are correctly “bridged” by axioms (46) and (47). As Table 4.1 shows, property (52) does not hold; this occurs because, when place `data_retrieved` is marked, the TA might not be in state `idle`.

$$\text{data_retrieved} \Rightarrow \Diamond_{(0, T]}(\text{try}) \tag{52}$$

Axiom (47) states that a `try` state is entered within $T/2$ if `data_retrieved` holds when `idle` holds. Then, a deeper analysis on the timing constraints suggests that this condition depends on the maximum transmission time T_3 of the TA, which defines the maximum delay between two consecutive occurrences of `idle`. If the system is in `data_retrieved` and not in `idle`, then the next `idle` state will be within T_3 instants in the future; moreover, `data_retrieved` will be unmarked within $T/2$. This suggests that the following property (53) is valid:

$$\square_{(0,T_3]}(\text{data_retrieved}) \Rightarrow \diamond_{(0,T]}(\text{try}) \quad (53)$$

This property also falls in the incompleteness region of the verification technique. However, the following slight relaxation of formula (53) can be proved to hold for the system:

$$\square_{(0,T_3+\delta]}(\text{data_retrieved}) \Rightarrow \diamond_{(0,T]}(\text{try}) \quad (54)$$

PR	T_1	T_2	T_3	T	K	PRE (min.)	CNF (hrs.)	SAT (hrs.)	N-VALID	# CL-10 ⁶
48: ϕ^+	3	6	18	30	90	1.9877	1.854	2.2322	\perp	12.4148
48: ϕ^-	3	6	18	30	90	3.0743	6.2533	5.3518	\perp	21.306
48: ϕ^+	3	9	36	30	90	2.425	2.5699	2.5368	\perp	12.7411
48: ϕ^-	3	9	36	30	90	3.3372	6.2202	5.0851	\perp	21.6323
48: ϕ^+	3	12	48	30	120	3.2059	3.6904	6.8226	\perp	17.2833
48: ϕ^-	3	12	48	30	120	4.5452	10.439	9.2688	\perp	29.117
49: ϕ^+	3	6	18	30	90	2.1074	1.9171	0.8101	T	12.8512
49: ϕ^-	3	6	18	30	90	3.1059	5.7381	3.017	T	21.7514
49: ϕ^+	3	9	36	30	90	2.6346	2.7726	0.9741	T	13.1775
49: ϕ^-	3	9	36	30	90	3.5125	6.4452	3.5557	T	22.0778
49: ϕ^+	3	12	48	30	120	3.6731	4.379	2.0955	T	17.8641
49: ϕ^-	3	12	48	30	120	5.1492	11.0093	5.0007	T	29.7098
50: ϕ^+	3	6	18	30	90	1.8887	1.7376	3.1524	\perp	12.0598
50: ϕ^-	3	6	18	30	90	2.9094	6.0154	3.427	T	20.931
50: ϕ^+	3	9	36	30	90	2.2002	2.3232	2.4845	\perp	12.3862
50: ϕ^-	3	9	36	30	90	3.1067	5.8341	4.6997	T	21.2573
50: ϕ^+	3	12	48	30	120	3.4446	4.1686	8.8680	\perp	16.8108
50: ϕ^-	3	12	48	30	120	4.1621	9.9533	13.1718	T	28.6179
51: ϕ^+	3	6	18	30	90	2.0715	1.6828	1.2976	T	12.1584
51: ϕ^-	3	6	18	30	90	3.0536	5.3665	3.9414	T	21.0296
51: ϕ^+	3	9	36	30	90	2.8152	2.2134	1.7645	T	12.4848
51: ϕ^-	3	9	36	30	90	3.7314	6.1665	3.6802	T	21.3559
51: ϕ^+	3	12	48	30	120	3.9268	4.5246	9.3435	\perp	16.9421
51: ϕ^-	3	12	48	30	120	4.8244	9.7484	14.8257	T	28.7491
50': ϕ^+	3	6	18	30	90	2.2399	2.3971	4.0335	T	12.8097
50': ϕ^-	3	6	18	30	90	3.3884	5.5905	4.5752	T	21.6645
50': ϕ^+	3	9	36	30	90	2.4788	2.2978	4.8259	T	13.136
50': ϕ^-	3	9	36	30	90	3.8369	7.3132	0.0036	T	21.9909
50': ϕ^+	3	12	48	30	120	4.7220	5.0607	13.3136	\perp	17.8088
50': ϕ^-	3	12	48	30	120	4.8557	9.7088	8.4951	T	29.5942
52: ϕ^+	3	6	12	30	75	1.5108	1.0502	0.4716	\perp	9.91056
52: ϕ^-	3	6	12	30	75	2.1418	3.1694	1.4723	\perp	17.3177
52: ϕ^+	3	3	15	30	75	1.5199	1.0564	0.4703	\perp	9.87584
52: ϕ^-	3	3	15	30	75	2.1586	3.1764	1.4473	\perp	17.2837
52: ϕ^+	3	6	18	30	75	1.5458	1.0706	0.5673	\perp	9.97865
52: ϕ^-	3	6	18	30	75	2.1978	3.2174	1.4323	\perp	17.3858
53: ϕ^+	3	6	12	30	75	1.6018	1.1108	0.8844	\perp	9.97312
53: ϕ^-	3	6	12	30	75	2.2909	3.3455	2.1095	T	17.3841
53: ϕ^+	3	3	15	30	75	1.6734	1.1945	0.6418	\perp	9.95542
53: ϕ^-	3	3	15	30	75	2.1638	3.2626	1.5792	T	17.3671
53: ϕ^+	3	6	18	30	75	1.7031	1.2210	0.9653	T	10.0752
53: ϕ^-	3	6	18	30	75	2.48	3.3642	1.1761	T	17.4862
54: ϕ^+	3	6	12	30	75	1.578	1.0879	1.2972	T	9.97879
54: ϕ^-	3	6	12	30	75	2.3035	3.2128	1.6002	T	17.3898
54: ϕ^+	3	3	15	30	75	1.6465	1.0986	0.7740	\perp	9.96109
54: ϕ^-	3	3	15	30	75	2.1604	3.1919	1.1408	T	17.3727
54: ϕ^+	3	6	18	30	75	1.6220	1.1249	0.8240	T	10.0809
54: ϕ^-	3	6	18	30	75	2.2892	3.2682	1.1178	T	17.4919

Table 2: Checking properties of the data monitoring system.

5 Discussion and Conclusion

In this paper we presented a technique to formally model and verify systems using different paradigms for different system parts. The technique hinges on MTL axiomatizations of the different modeling notations, which provide a common formal ground for the various modeling languages, on which fully-automated verification techniques are built. We provided an MTL axiomatization of a subset of TPN, a typical asynchronous operational formalism, and showed how models could be built by formally combining together TPN and TA (a classic synchronous operational notation, for which an axiomatization has been provided in [12]). In addition, we showed how the approach allows users to integrate in the same model parts described through a continuous notion of time, and parts described through a discrete notion of time.

Practical verification of systems modeled through the multi-paradigm approach is possible through the \mathcal{Z} ot bounded satisfiability checker, for which plugins supporting the various axiomatized notations have been built.

The technique has been validated on a non trivial example of data monitoring system. The experimental results show the feasibility of the approach, through which we have been able to investigate the validity (or, in some cases, the non validity) of some properties of the system. As described in Section 4, the verification phase has provided useful insights on the mechanisms and on the timing features of the modeled system, which led us to re-evaluate some of our initial beliefs on the system properties.

It is clear from our experiments that, unsurprisingly, the technique suffers from two main drawbacks: the incompleteness of the verification approach by discretization evidenced in [11], which prevented us, in some cases, to get conclusive answers on some analyzed properties; and the computational complexity of our method, which is based on the direct translation of TPN and TA into MTL formulas, approximated into discrete ones, and then encoded into SAT. This makes proofs considerably lengthier as the size of the domains, and especially of the temporal one, increases, as evidenced by Table 4.1. Nevertheless, we maintain that the results we obtained are promising, and show the applicability of the technique on non trivial systems. This claim is supported on the one hand by the sophistication of the properties we have been able to prove (or disprove): it is inevitable that verification over continuous real-time has a high computational cost. On the other hand, while incompleteness is a hurdle to the full applicability of the technique, in practice it can be mitigated quite well, usually by slightly relaxing the real-time timing requirements under verification in a way that does not usually alter the gist of what is being verified.

In our future research on this topic we plan to address the two main drawbacks evidenced above. First, we will work on extending the verification technique to expand its range of applicability and reduce its region of incompleteness. Also, we will study more efficient implementations for the \mathcal{Z} ot plugins through which the various modeling notations are added to the framework: we believe that more direct (therefore more compact, both in the literals and clause numbers) encodings into SAT of the TPN and TA axiomatizations should significantly improve the efficiency of the tool.

In particular, we have not yet tackled the problem of optimizing the encodings of the TPN and TA axiomatizations into the SAT problem. We expect that significant improvements on the duration of the proofs can be gained through optimized encodings that reduce, on the one hand, the time needed to put formulas in the conjunctive normal form that is required as input by SAT solvers, and, on the other hand, the number of literals required to represent TPN and TA as SAT problems.

References

- [1] R. Alur and D. L. Dill. A theory of timed automata. *Theor. Comp. Sci.*, 126(2):183–235, 1994.
- [2] R. Alur, T. Feder, and T. A. Henzinger. The benefits of relaxing punctuality. *Journal of the ACM*, 43(1):116–146, 1996.
- [3] R. Alur and T. A. Henzinger. Logics and models of real time: A survey. In *Proc. of Real-Time: Theory in Practice*, volume 600 of *LNCS*, pages 74–106, 1992.
- [4] R. Alur and T. A. Henzinger. Real-time logics: Complexity and expressiveness. *Information and Computation*, 104(1):35–77, 1993.
- [5] F. Cassez and O. H. Roux. Structural translation from time Petri nets to timed automata. *Journal of Systems and Software*, 79(10):1456–1468, 2006.
- [6] A. Cerone and A. Maggiolo-Schettini. Time-based expressivity of time Petri nets for system specification. *Theor. Comp. Sci.*, 216(1–2):1–53, 1999.
- [7] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.
- [8] M. Felder, D. Mandrioli, and A. Morzenti. Proving properties of real-time systems through logical specifications and Petri net models. *IEEE Trans. on Soft. Eng.*, 20(2):127–141, 1994.
- [9] C. A. Furia, D. Mandrioli, A. Morzenti, and M. Rossi. Modeling time in computing. Technical Report 2007.22, DEI, Politecnico di Milano, January 2007.
- [10] C. A. Furia, D. Mandrioli, A. Morzenti, and M. Rossi. Modeling time in computing: a taxonomy and a comparative survey. *ACM Computing Surveys*, to appear. Also available as <http://arxiv.org/abs/0807.4132>.
- [11] C. A. Furia, M. Pradella, and M. Rossi. Automated verification of dense-time MTL specifications via discrete-time approximation. In *Proc. of FM'08*, volume 5014 of *LNCS*, pages 132–147, 2008.
- [12] C. A. Furia, M. Pradella, and M. Rossi. Practical automated partial verification of multi-paradigm real-time models. In *Proc. of ICFEM'08*, volume 5256/-1 of *LNCS*, pages 298–317, 2008.
- [13] C. A. Furia and M. Rossi. Integrating discrete- and continuous-time metric temporal logics through sampling. In *Proc. of FORMATS'06*, volume 4202 of *LNCS*, pages 215–229, 2006.
- [14] C. Heitmeyer and D. Mandrioli, editors. *Formal Methods for Real-Time Computing*. John Wiley & Sons, 1996.
- [15] R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
- [16] M. Pradella. Zot. <http://home.dei.polimi.it/pradella>, March 2007.

- [17] M. Pradella, A. Morzenti, and P. San Pietro. The symmetry of the past and of the future: bi-infinite time in the verification of temporal properties. In *Proc. of ESEC/FSE 2007*, 2007.
- [18] OMG Unified Modeling Language (OMG UML) Superstructure, v2.2. Technical Report formal/2009-02-02, Object Management Group, 2009.
- [19] M. von der Beeck. A comparison of statecharts variants. In *Proc. of FTRTFT*, volume 863 of *LNCS*, pages 128–148, 1994.