Half-order Modal Logic: How To Prove Real-time Properties*

Thomas A. Henzinger†
Department of Computer Science
Stanford University

Abstract. We introduce a novel extension of propositional modal logic that is interpreted over Kripke structures in which a value is associated with every possible world. These values are, however, not treated as full first-order objects; they can be accessed only by a very restricted form of quantification: the "*freeze*" quantifier binds a variable to the value of the current world. We present a complete proof system for this ("*half-order*") modal logic.

As a special case, we obtain the real-time temporal logic TPTL of [AH89]: the models are restricted to infinite sequences of states, whose values are monotonically increasing natural numbers. The ordering relation between states is interpreted as temporal precedence, while the value associated with a state is interpreted as its "real" time. We extend our proof system to be complete for TPTL, and demonstrate how it can be used to derive real-time properties.

1 Introduction

Our main result is the presentation of a complete proof system for the real-time temporal logic TPTL of [AH89]. This constitutes an important step in the development of a formal verification methodology for real-time systems.

Temporal logic has been established as a suitable formalism for the specification and verification of a large class of properties of reactive systems (see, for example, [MP89] for a comprehensive presentation of the temporal methodology). One shortcoming of propositional temporal logic is that it admits only the treatment of *qualitative* time requirements, which contain temporal ordering as the only information about the times of events. A typical example is the requirement that every request p is, "eventually," followed by a response q:

$$\Box(p \rightarrow \Diamond q).$$

Because temporal logic abstracts over the actual times at which events occur, it cannot express *quantitative* time requirements, such as every request p being followed by a response q within 5 time

*A preliminary version of this paper appeared in the *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing* (PODC 1990), pp. 281–296.

†This research was supported by an IBM graduate fellowship, by the NSF grant CCR-8812595, by the DARPA contract N00039-84-C-0211, and by the USAF Office of Scientific Research under contracts 88-0281 and 90-0057.

units. This is inadequate for systems whose correctness depends crucially on such "hard" real-time properties, as does the correctness of many communication protocols and control circuits.

That is why researchers have been looking for languages that are interpreted over state structures in which a time is associated with every state. The obvious solution is to employ a first-order temporal logic, with a state variable T that represents, in every state, the current value of the time (e.g., [PH88], [Os89]); the real-time response property stated above can then be written as

$$\Box\forall x.((p \wedge x = T) \rightarrow \Diamond(q \wedge T \leq x + 5)),$$

where the rigid variable x is used to record the time of any request p.

We have argued that this upgrade from a propositional logic to its full first-order version may be both unnecessary and expensive: only a very restricted form of quantification achieves the expressive power of first-order linear temporal logic, while retaining a much simpler decision problem ([AH89], [AH90]). Furthermore, the sacrifice of ordinary universal and existential quantification in favor of the single *freeze*[1] quantifier "$x.$", which binds ("freezes") the variable x to the (unique) value of the time in the current state (and therefore is its own dual), yields a natural specification language:

$$\Box x.(p \rightarrow \Diamond y.(q \wedge y \leq x + 5)).$$

(Read this formula as "Whenever there is a request p, and the variable x is frozen to the current time, the request is followed by a response q, at time y, such that y is at most $x + 5$.")

Thus we have made a convincing case for the use of a small fragment of first-order linear temporal logic, *timed propositional temporal logic* (TPTL). We also showed that TPTL yields to tableau-based verification techniques ([AH89]). A complete formal approach to real-time specification and verification should, however, encompass both syntactic and semantic methods. Here we develop a complete proof system for TPTL, which complements the model-checking algorithm given in [AH89]. This is of particular importance for real-time systems that cannot be represented as finite state graphs.

Just as proof systems for propositional temporal logic consist of a general, modal, part and special axioms for linear structures, we arrive at the proof system for TPTL in two steps. First (in Section 2), we axiomatize the freeze quantifier for arbitrary modal logics that are interpreted over Kripke structures in which a value is associated with every possible world, completely independent of the notion of "time." Since these modal logics are fragments of the corresponding first-order versions, we dub them *half-order*. We also show that half-order modal logic generalizes the classical first-order predicate calculus, which can be embedded.

Secondly (in Section 3), we obtain half-order *temporal* logic by restricting ourselves to certain linear models. We interpret possible worlds as system states, the accessibility relation as temporal ordering between states, and the value associated with a state as its time. The resulting *timed state sequences* are precisely the models of TPTL. By adding appropriate axioms for linear structures and the timing constraints admitted in TPTL, we obtain a complete proof system for TPTL. In fact, the choice of timing constraints of TPTL turns out to be crucial; we show that half-order temporal logic in general is Π^1_1-hard, and hence not axiomatizable.

In the final section, we indicate several other possible applications of the freeze quantifier, both in dynamic logic and modal logics of knowledge.

[1] This terminology was suggested by Amir Pnueli.

2 Half-order Modal Logic

We introduce modal logics that are interpreted over Kripke structures each of whose possible worlds (states) s has a value $|s|$ associated with it. Ordinary first-order function and predicate symbols, including equality, perform operations and tests on these values. However, instead of ordinary universal (and existential) quantification, the access to values is kept extremely local: the "*freeze*" quantifier "$x.$" binds x to the value that is associated with the current state.

For example, the formula $x.\Diamond y.\, p(x,y)$ is true in a model with initial state s iff there is a state t accessible from s such that the relation p, as interpreted in t, holds between the value $|s|$ associated with s and the value $|t|$ associated with t.

There is a wide, and largely confusing, variety of different ways to add conventional quantification to modal logic, for only some of which completeness results have been achieved, and some of which are known to be incomplete (see [Ga84] for an excellent survey). The situation for the freeze quantifier is, fortunately, much cleaner: we show that modal logics with freeze quantification are axiomatizable, yet not necessarily decidable, fragments of certain corresponding first-order modal logics. This explains the attribute "*half-order*" for the freeze quantifier.

2.1 Syntax and semantics

The formulas of half-order modal logic are built from first-order atoms, which include equations, by propositional connectives, the modal operator $\Box$, and the freeze quantifier.

Let V be an infinite set of variables, and F and P be sets of function and predicate symbols, respectively. We assume that all of these sets can be effectively enumerated. The *terms* π, *atomic formulas* α, and *formulas* ϕ of half-order modal logic are inductively defined as follows:

$$\begin{aligned}
\pi &:= x \mid f\bar{\pi} \\
\alpha &:= \pi_1 = \pi_2 \mid p\bar{\pi} \\
\phi &:= \alpha \mid \mathbf{false} \mid \phi_1 \rightarrow \phi_2 \mid \Box\phi \mid x.\,\phi
\end{aligned}$$

for $x \in V$, $f \in F$, and $p \in P$. (We write $\bar{\pi}$ to denote a tuple of terms; for example, if f is a binary function symbol then $f\bar{\pi}$ stands for $f\pi_1\pi_2$, or $f(\pi_1,\pi_2)$.)

Propositional connectives such as $\neg$, $\wedge$, and $\vee$ are defined in terms of $\rightarrow$ and **false** as usual; $\Diamond\phi$ is an abbreviation for $\neg\Box\neg\phi$. Throughout the paper, we use π and α (possibly subscripted) to stand for arbitrary terms and atomic formulas, respectively, and ϕ, ψ, and φ to denote formulas.

An *interpretation*

$$\mathcal{M} = (\mathcal{S}, \rightarrow_\Box, \mathcal{U}, ||, [\![x]\!]_{x\in V}, [\![f]\!]_{f\in F}, [\![p]\!]_{p\in P}, s_0)$$

for half-order modal logic consists of

- a set $\mathcal{S}$ of states,
- an accessibility relation $\rightarrow_\Box \subseteq \mathcal{S}^2$ on the states,
- a set $\mathcal{U}$ of values,
- a value function $||: \mathcal{S} \rightarrow \mathcal{U}$ that associates a value $|s|$ with every state s,
- a rigid assignment function $[\![x]\!] \in \mathcal{U}$ for all variables $x \in V$,
- a rigid assignment function $[\![f]\!]: \bar{\mathcal{U}} \rightarrow \mathcal{U}$ for all function symbols $f \in F$,
- a flexible assignment function $[\![p]\!]: \mathcal{S} \rightarrow 2^{\bar{\mathcal{U}}}$ for all predicate symbols $p \in P$, and
- an initial state $s_0 \in \mathcal{S}$.

Note that all terms are given a state-independent (rigid) meaning, while the interpretation of predicate symbols is state-dependent (flexible). The rigidity restriction on terms is required for

the completeness proof; the flexibility of predicate symbols is necessary to cover TPTL, whose propositions are interpreted in a state-dependent fashion.

The interpretation $\mathcal{M}$ is a *model* of the formula ϕ iff $\mathcal{M} \models \phi$, for the following inductive definition of the truth predicate $\models$:

$$\begin{array}{lll}
\mathcal{M} \models \pi_1 = \pi_2 & \text{iff} & [\![\pi_1]\!] = [\![\pi_2]\!], \\
 & & \text{for } [\![f\bar{\pi}]\!] = [\![f]\!]([\![\bar{\pi}]\!]) \\
\mathcal{M} \models p\bar{\pi} & \text{iff} & [\![\bar{\pi}]\!] \in [\![p]\!](s_0) \\
\mathcal{M} \not\models \mathbf{false} & & \\
\mathcal{M} \models \phi_1 \rightarrow \phi_2 & \text{iff} & \mathcal{M} \models \phi_1 \text{ implies } \mathcal{M} \models \phi_2 \\
\mathcal{M} \models \Box\phi & \text{iff} & \mathcal{M}[s_0 := t] \models \phi \\
 & & \text{for all } t \in \mathcal{S} \text{ with } s_0 \rightarrow_\Box t \\
\mathcal{M} \models x.\phi & \text{iff} & \mathcal{M}[[\![x]\!] := |s_0|] \models \phi.
\end{array}$$

Here $\mathcal{M}[s_0 := t]$ denotes the interpretation that differs from $\mathcal{M}$ only in its initial state, t; the interpretation $\mathcal{M}[[\![x]\!] := |s_0|]$ differs from $\mathcal{M}$ only in its assignment function for x. Thus, the semantic clause for formulas of the form $\Box\phi$ specifies that ϕ is interpreted in all states accessible from the current state. The clause for $x.\phi$ asserts that all occurrences of x in ϕ refer to the value that is associated with the current state.

The formula ϕ is *satisfiable* (*valid*) iff some (every) interpretation is a model of ϕ. Two formulas are *equivalent* iff they have the same models.

Observe that we can faithfully embed half-order modal logic into a first-order modal logic with the set $\mathcal{U}$ of values as constant domain for all states, and rigid terms, with the exception of *one* flexible constant symbol, say T, that denotes, in every state, the value associated with that state (i.e., $[\![T]\!](s) = |s|$). The freeze quantifier $x.\phi$ is translated as

$$\forall x.(x = T \rightarrow \phi)$$

(or, equivalently, $\exists x.(x = T \wedge \phi)$). In other words, half-order modal logic captures the fragment of first-order modal logic in which every rigid (global) variable is, immediately upon introduction, bound to the current value of the flexible (state) variable T. This makes the conventional first-order quantifiers superfluous.

2.2 Proof system

Let K_P be the propositional modal logic that is determined by the class of all Kripke structures. We extend the proof system for K_P by axioms and inference rules for both the freeze quantifier and equality.

Recall that the deductive calculus for K_P consists of a complete proof system **PROP** for propositional logic, say,

PROP1 all tautologies are axioms
PROP2 from ϕ_1 and $\phi_1 \rightarrow \phi_2$ infer ϕ_2,

as well as the following axiom schema and necessitation rule, which completely characterize the modal operator $\Box$ with respect to Kripke semantics:

K1 $\Box(\phi_1 \rightarrow \phi_2) \rightarrow \Box\phi_1 \rightarrow \Box\phi_2$
K2 from ϕ infer $\Box\phi$.

Let us abbreviate $\phi \rightarrow \Box\psi$ to $\phi \Rightarrow \psi$ ($\Rightarrow$ associates to the right, as does $\rightarrow$). The freeze quantifier is characterized by an axiom schema that asserts its functionality,

Q1 $\quad x.(\phi_1 \rightarrow \phi_2) \leftrightarrow (x.\phi_1 \rightarrow x.\phi_2)$,

and an introduction rule for every modal context,

Q2 from $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow \psi$
infer $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow x.\psi$,
provided that $x \notin \phi_i$ for all $1 \leq i \leq n$

(we write $x \in \phi$ iff the variable x occurs freely in ϕ). The simplest instances of this rule (take $n = 0$) are of the form

Q2* $\quad$ from ϕ infer $x.\phi$.

Only the elimination of vacuous quantifier occurrences is sound:

Q3 $\quad x.\phi \leftrightarrow \phi$ if $x \notin \phi$.

Let $\phi\langle\pi_1 := \pi_2\rangle$ ($\phi[\pi_1 := \pi_2]$) denote a formula that results from ϕ by safely replacing zero, one, or more (all, respectively) free occurrences of π_1 by π_2. *Safe* replacement means, as usual, that no free occurrence of π_1 that is replaced, is within the scope of a quantifier binding a variable of π_2; whenever we write $\phi[\pi_1 := \pi_2]$, there is the implicit condition that *all* free occurrences of π_1 in ϕ can be safely replaced by π_2.

We add conventional congruence axioms for equality, for instance,

EQ1 $\quad \pi = \pi$
EQ2 $\quad \pi_1 = \pi_2 \rightarrow \alpha \rightarrow \alpha\langle\pi_1 := \pi_2\rangle$,

two axiom schemata that assert the rigidity of terms,

RIG1 $\quad \pi_1 = \pi_2 \rightarrow \Box\pi_1 = \pi_2$
RIG2 $\quad \pi_1 \neq \pi_2 \rightarrow \Box\pi_1 \neq \pi_2$,

and an axiom schema that states that the value associated with every state is unique,

QEQ $\quad x.y.x = y$.

To summarize, we are given the *logical axioms* **PROP1**, **K1**, **Q1**, **Q3**, **EQ1-2**, **RIG1-2**, and **QEQ**, and the *inference rules* **PROP2**, **K2**, and **Q2**. A half-order *normal logic* is a set of formulas that contains all logical axioms and is closed under all inference rules.

Note that, since the proof system includes **PROP** and **K1-2**, every normal logic contains all instances of valid schemata of the propositional modal logic K_P. We demonstrate the use of the proof system by deriving some additional theorems of half-order modal logic that will be useful later.

LEMMA [**Sample theorems**]. *Any normal logic contains the following formulas:*

Q4 $\quad x.\neg\phi \leftrightarrow \neg x.\phi$
EQ3 $\quad \pi_1 = \pi_2 \rightarrow \phi \rightarrow \phi\langle\pi_1 := \pi_2\rangle$
EQ4 $\quad x.x = \pi \rightarrow (x.\phi \leftrightarrow \phi[x := \pi])$
VAR1 $\quad x.\phi \rightarrow y.\phi[x := y]$ *if* $y \notin \phi$
VAR2 $\quad \phi \leftrightarrow \phi'$ *if* ϕ' *results from* ϕ *by safe renaming of bound variables*
VAR3 $\quad x.y.\phi \leftrightarrow x.\phi[y := x]$. ∎

PROOF. **Q4** states that the freeze quantifier is its own dual. It follows from **Q1**:

$$x.(\phi_1 \rightarrow \mathbf{false}) \leftrightarrow (x.\phi_1 \rightarrow x.\mathbf{false})$$

by **Q3** and **PROP**.

EQ3 generalizes the equality axiom **EQ2** to arbitrary formulas ϕ. To establish it, we use induction on the structure of ϕ. The atomic base case holds by **EQ2**. The propositional cases follow from the induction hypothesis by **PROP**. Observe that

$$\pi_1 = \pi_2 \rightarrow \Box\phi \rightarrow (\Box\phi)\langle\pi_1 := \pi_2\rangle$$

equals

$$\pi_1 = \pi_2 \rightarrow \Box\phi \rightarrow \Box\phi\langle\pi_1 := \pi_2\rangle,$$

which follows from the induction hypothesis by **K2**, **K1**, **RIG1**, and **PROP**. If $x \notin \pi_1$ and $x \notin \pi_2$, then $(x.\phi)\langle\pi_1 := \pi_2\rangle$ equals $x.\phi\langle\pi_1 := \pi_2\rangle$, and

$$\pi_1 = \pi_2 \rightarrow x.\phi \rightarrow x.\phi\langle\pi_1 := \pi_2\rangle,$$

follows from the induction hypothesis by **Q2***, **Q1**, **Q3**, and **PROP**. If, on the other hand, $x \in \pi_1$ or $x \in \pi_2$, then $(x.\phi)\langle\pi_1 := \pi_2\rangle$ must be $x.\phi$; the corresponding inductive step holds by **PROP**.

From now on, we will often omit to mention applications of **PROP** explicitly. Similarly to our use of **PROP**, we may refer to any of **K1-2** and **Q1-4** simply by **K** or **Q**. We write **Q*** if all applications of the rule **Q2** are instances of **Q2***.

EQ4 will be used extensively in the completeness proof, and can be derived from **EQ3** by **Q***.

VAR1 shows that bound variables can be renamed at the top level of formulas; it follows from **EQ4** by **Q*** and **QEQ**. **VAR2** generalizes **VAR1** and is shown by structural induction.

VAR3 demonstrates that adjacent quantifiers can be combined; it follows from **EQ4** by **Q*** and **QEQ**. Together with **Q1**, **VAR3** implies that every formula ϕ is equivalent to some formula $x.\phi'$ that contains at most one quantifier per modal level; that is, every quantifier in ϕ' *follows* a modal operator. Alternatively, any formula can be put into a normal form in which every quantifier *precedes* a modal operator or an atomic formula. ∎

Let Φ be a set of formulas and Λ_Φ the intersection of all normal logics containing Φ. Clearly, Λ_Φ is again a normal logic; the formulas of Φ are called the *nonlogical axioms* of Λ_Φ. In particular, $\mathsf{K} = \Lambda_\emptyset$ is the smallest normal logic, and $\perp = \Lambda_{\{\mathbf{false}\}}$, the set of all formulas, is the largest one.

Our goal is to show that K is precisely the set of all valid formulas, that is, that the given proof system is both sound and complete for half-order modal logic. First we show that every formula of K is true under every interpretation; the completeness proof is deferred to the next section.

LEMMA [**Soundness**]. *The logical axioms are valid. If all of the antecedents of an inference rule are valid, then so is the consequent.* ∎

PROOF. The soundness of **PROP**, **K1**, **Q1**, **EQ1-2**, **RIG1-2**, and **QEQ** follows immediately from the definition of truth under an interpretation. The argument for **K2** is the same as in propositional modal logic.

To show **Q3** to be sound, we use the following fact:

$$\begin{array}{l}\text{for all } x \notin \phi \text{ and values } u \text{ of } \mathcal{M},\\ \mathcal{M} \models \phi \;\text{ iff }\; \mathcal{M}[[\![x]\!] := u] \models \phi.\end{array} \qquad (\dagger)$$

This can be established by induction on the structure of ϕ.

It remains to be shown that **Q2** is sound. Suppose that $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow \psi$ is valid and $x \notin \phi_i$ for all $1 \leq i \leq n$. Now consider an arbitrary interpretation $\mathcal{M}$, and show that $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow x.\psi$ is true under $\mathcal{M}$. Let s_0 be the initial state of $\mathcal{M}$, and

$$s_0 = t_1 \rightarrow_\Box \cdots \rightarrow_\Box t_{n+1}.$$

Assume that $\mathcal{M}[s_0 := t_i] \models \phi_i$ for all $1 \leq i \leq n$, and show that $x.\psi$ is true under $\mathcal{M}[s_0 := t_{n+1}]$.

Since $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow \psi$ is valid, it is in particular true under $\mathcal{M}[[\![x]\!] := |t_{n+1}|]$. By (†), we infer from $\mathcal{M}[s_0 := t_i] \models \phi_i$ and $x \notin \phi_i$ that $\mathcal{M}[s_0 := t_i, [\![x]\!] := |t_{n+1}|] \models \phi_i$ for all $1 \leq i \leq n$. The desired conclusion $\mathcal{M}[s_0 := t_{n+1}, [\![x]\!] := |t_{n+1}|] \models \psi$ follows. ∎

2.3 Completeness

Our completeness proof is typical for quantified modal logics, and combines techniques from both propositional modal logic and classical first-order logic. The organization of the proof follows largely similar proofs given in [Ga84].

As is common in propositional modal logic, we construct, for any consistent formula, a model whose states are maximally consistent sets. The basic idea of the construction of this "canonical" model is to guarantee that all formulas that are contained in a state, are true in that state.

As usual in Henkin-type proofs of the completeness of classical first-order logic, we take as the set of values the equivalence classes of terms under equality (which requires the rigidity of terms). A completeness condition has to be put on the maximally consistent sets to assure that all of them have values associated with them. So in order to give the canonical-model construction, we have to develop the notions of consistency and completeness of sets of formulas first.

Given a normal logic Λ, we write $\Phi \vdash_\Lambda \phi$ iff

$$\phi_1 \rightarrow \cdots \rightarrow \phi_n \rightarrow \phi \in \Lambda$$

for some finite subset $\{\phi_i \mid 1 \leq i \leq n\}$ of Φ, and $\vdash_\Lambda \phi$ iff $\phi \in \Lambda$. Throughout this section, we assume $\Lambda \neq \bot$ to be fixed, and suppress the subscript of the derivability predicate $\vdash$.

The concept of consistency is the familiar one. A set Φ of formulas is called *consistent* iff $\Phi \nvdash$ **false**; it is *maximally* consistent iff, furthermore, either $\phi \in \Phi$ or $\neg\phi \in \Phi$ for all formulas ϕ. The following lemma is, as usual, established using **PROP**.

LEMMA [**Consistency**].

(i) $\Phi \cup \{\neg\phi\}$ *is consistent iff* $\Phi \nvdash \phi$.
(ii) If Φ *is consistent, then either* $\Phi \cup \{\phi\}$ *or* $\Phi \cup \{\neg\phi\}$ *is consistent.*
(iii) For any maximally consistent set Φ, $\phi_1 \rightarrow \phi_2 \in \Phi$ *iff* $\phi_1 \notin \Phi$ *or* $\phi_2 \in \Phi$. ∎

By $\phi - x$ we denote a *set* of formulas that result from binding all free occurrences of x in ϕ by a single quantifier. More precisely, $\phi - x$ is the smallest set satisfying the following condition: if ϕ is of the form

$$\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow \psi,$$

where $x \notin \phi_i$ for all $1 \leq i \leq n$, then

$$\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow x.\psi \in \phi - x.$$

Note that, in particular, $x.\phi \in \phi - x$.

A set Φ of formulas is *complete* iff $\Phi \vdash \phi[x := \pi]$ for all terms π implies $\Phi \vdash \phi'$ for all $\phi' \in \phi - x$. By part (i) of the *consistency* lemma, it follows that completeness is equivalent to the condition that, if $\Phi \cup \{\neg\phi'\}$ is consistent and $\phi' \in \phi - x$ for some formula ϕ and variable x, then there is a term π such that $\Phi \cup \{\neg\phi[x := \pi]\}$ is consistent.

In particular, the completeness of a set Φ ensures that, whenever Φ contains the formulas $\phi[x := \pi]$ for all terms π, then Φ entails $x.\phi$. This means, intuitively speaking, that some term π is interpreted as the value of the current state. The general form of the completeness condition is necessary to guarantee this property, which allows us to assign terms as values to complete sets, for all states in the canonical model.

Maximally consistent, complete sets are called *saturated*. Given a consistent formula ϕ_0, we will define a model for ϕ_0 whose states are saturated sets. Part (i) of the following lemma ensures that ϕ_0 is contained in some saturated set, while part (ii) guarantees that that there are enough such sets for constructing a model. Note that the general form of the quantifier introduction rule **Q2** is needed to show the former.

LEMMA [**Existence of saturated extensions**].

(i) Every finite consistent set has a saturated extension.
(ii) Every complete consistent set has a saturated extension. ∎

PROOF. Both parts are shown by applying variants of the Lindenbaum procedure to extend a consistent set Φ_0 to a maximally consistent set Φ. Enumerate all formulas $\phi_1, \phi_2, \ldots$, and let Φ_{i+1}, for all $i \geq 0$, be either $\Phi \cup \{\phi_i\}$, if this set is still consistent, or $\Phi_i \cup \{\neg\phi_i\}$, otherwise. Each set Φ_i is consistent by part (ii) of the *consistency* lemma, implying that $\phi = \bigcup_{i\geq 0} \Phi_i$ is maximally consistent.

(i) Suppose that Φ_0 is finite. Then we can guarantee the completeness of Φ as follows: whenever Φ_i is extended by a formula $\neg\phi'$ and $\phi' \in \phi - x$ for some formula ϕ and variable x, then we add also $\neg\phi[x := y]$, for some new variable $y \notin \Phi_i \cup \{\phi\}$. Although at every stage several new formulas may be added, each set Φ_i is finite, which assures the existence of new variables at all future stages.

It remains to be shown that this process preserves consistency. Assume that $\Phi \cup \{\neg\phi', \neg\phi[x := y]\}$ is inconsistent, that is,

$$\vdash \Phi_f \rightarrow \neg\phi' \rightarrow \neg\phi[x := y] \rightarrow \mathbf{false} \tag{1}$$

for some conjunction Φ_f of formulas in Φ. We derive a contradiction, by showing that, in this case, already

$$\vdash \Phi_f \rightarrow \neg\phi' \rightarrow \mathbf{false}, \tag{2}$$

implying the inconsistency of $\Phi \cup \{\neg\phi'\}$. Let ϕ' be $\phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow x.\psi$, with $x \notin \phi_i$ for all $1 \leq i \leq n$. Use **Q** to infer from (1) that

$$\vdash \Phi_f \rightarrow \neg\phi' \rightarrow \phi_1 \Rightarrow \cdots \Rightarrow \phi_n \Rightarrow y.\psi[x := y],$$

and conclude (2) by **VAR2** and **PROP**.

(ii) First we show that, if a set Φ is complete, then so is $\Phi \cup \{\phi\}$. Assume that $\Phi \cup \{\phi\} \vdash \varphi[x := \pi]$ for all terms π, let x be such that $x \notin \phi$ and $\varphi' \in \varphi - x$; by **VAR2** it suffices to show that $\Phi \cup \{\phi\} \vdash \varphi'$. From our assumption, it follows by **PROP** that $\Phi \vdash \phi \rightarrow \varphi[x := \pi]$ for all terms π; hence $\Phi \vdash \phi \rightarrow \varphi'$ by the completeness of Φ (and $\mathbf{Q}^*$, in case φ' equals $x.\varphi$). The desired conclusion follows.

Now suppose that Φ_0 is complete. Then all the finite extensions Φ_i, $i \geq 0$, are complete; that is, whenever Φ_i is about to be extended by the formula $\neg\phi'$ and $\phi' \in \phi - x$ for some formula ϕ, there exists a term π such that $\Phi_i \cup \{\neg\phi[x := \pi]\}$ is consistent. We extend Φ_i in this fashion, and continue the Lindenbaum process by checking whether $\neg\phi'$ can still be added consistently. Since every formula contains only a finite number of quantifiers, each stage is completed within a finite number of steps. ■

In the canonical model, the value associated with a state (saturated set) s containing $x.x = \pi$ will be the equivalence class of the term π under equality. The next lemma guarantees that such a term exists (part (iv)), and that equality is a congruence relation (parts (i) to (iii)).

LEMMA [**Value of saturated sets**]. *For every saturated set s:*

(i) $\approx_s = \{(\pi_1, \pi_2) \mid \pi_1 = \pi_2 \in s\}$ *is an equivalence relation.*
(ii) $\bar{\pi}_1 \approx_s \bar{\pi}_2$ *implies that* $f\bar{\pi}_1 \approx_s f\bar{\pi}_2$ *for all* $f \in F$.
(iii) $\bar{\pi}_1 \approx_s \bar{\pi}_2$ *and* $p\bar{\pi}_1 \in s$ *implies* $p\bar{\pi}_2 \in s$ *for all* $p \in P$.
(iv) $\{\pi \mid x.x = \pi \in s\} = [\pi]_{\approx_s}$ *for some term* π. ■

PROOF. Parts (i), (ii), and (iii) follow from **EQ**.
(iv) **QEQ** and **Q4** imply that $\neg y.\neg x.\, y = x \in s$. Since s is complete, $x.\pi_s = x \in s$ for some term π_s not containing x; thus $x.x = \pi_s \in s$ by **EQ** and $\mathbf{Q}^*$. From **EQ4**:

$$x.x = \pi_1 \rightarrow (x.x = \pi_2 \leftrightarrow \pi_1 = \pi_2),$$

we infer that $[\pi_s]_{\approx_s} = \{\pi \mid x.x = \pi \in s\}$. ■

Now we are ready to define the *canonical model* for a given consistent formula ϕ_0. Let

$$\mathcal{M}(\phi_0) = (\mathcal{S}, \rightarrow_\Box, \mathcal{U}, |\,|, [\![x]\!]_{x \in V}, [\![f]\!]_{f \in F}, [\![p]\!]_{p \in P}, s_0)$$

be an interpretation such that:

- s_0 is a saturated extension of $\{\phi_0\}$. By $\approx$ we denote $\approx_{s_0}$.
- $\mathcal{S}$ is the set of saturated sets s with $\approx_s = \approx$.
- $s \rightarrow_\Box t$ iff $\phi \in t$ for all $\Box\phi \in s$.
- $\mathcal{U}$ is the set of equivalence classes $[\pi]_\approx$.
- $|s| = \{\pi \mid x.x = \pi \in s\}$.
- $[\![x]\!] = [x]_\approx$.
- $[\![f]\!]([\bar{\pi}]_\approx) = [f\bar{\pi}]_\approx$.
- $[\bar{\pi}]_\approx \in [\![p]\!](s)$ iff $p\bar{\pi} \in s$.

Note that the interpretation $\mathcal{M}(\phi_0)$ is well-defined: a saturated extension of $\{\phi_0\}$ exists by part (i) of the lemma on the *existence of saturated extensions*; the lemma on the *value of saturated sets* guarantees that $\approx$ is an equivalence relation, that $|s| \in \mathcal{U}$, and that $[\![f]\!]$ and $[\![p]\!]$ are properly defined.

The values of the canonical model are equivalence classes of terms under equality. It is straightforward to show, by structural induction, that all terms are interpreted as themselves, modulo equality.

LEMMA [**Term model**]. $[\![\pi]\!] = [\pi]_\approx$. ■

The states of the canonical model are saturated sets. The following main theorem shows that every state s of $\mathcal{M}(\phi_0)$ contains precisely the formulas that are true at s. Since $\phi_0 \in s_0$, it follows immediately that the interpretation $\mathcal{M}(\phi_0)$ is a model of ϕ_0.

THEOREM [**Canonical model**]. *For all* $s \in \mathcal{S}$, $\phi \in s$ *iff* $\mathcal{M}(\phi_0)[s_0 := s] \models \phi$. ■

PROOF. We apply induction on the structure of ϕ. The atomic cases follow from the *term-model* lemma. The propositional cases are consequences of the induction hypothesis and part (iii) of the *consistency* lemma.

If $\Box\phi \in s$ and $s \rightarrow_\Box t$, then $\phi \in t$ and, by the induction hypothesis, $\mathcal{M}[s_0 := t] \models \phi$. Now assume that $\Box\phi \notin s$, that is, $\neg\Box\phi \in s$, and show that $\mathcal{M}[s_0 := s] \not\models \Box\phi$. This follows from the *succession* lemma (see below) by the induction hypothesis.

For the quantifier case, choose a term π_s such that $x.x = \pi_s \in s$. Consequently, $\mathcal{M}[s_0 := s] \models x.\phi$ iff $\mathcal{M}[s_0 := s, [\![x]\!] := |s_0|] \models \phi$ iff, by the *substitution* lemma (below), $\mathcal{M}[s_0 := s] \models \phi[x := \pi_s]$ iff, by the induction hypothesis, $\phi[x := \pi_s] \in s$. From **EQ4**:

$$x.x = \pi_s \rightarrow (x.\phi \leftrightarrow \phi[x := \pi_s]),$$

we conclude that $\mathcal{M}[s_0 := s] \models x.\phi$ iff $x.\phi \in s$. ■

In the proof of the *canonical-model* theorem, we invoked the following two lemmas. To show that substitution behaves as expected, can be done by straightforward structural induction. To show that there are enough states in the canonical model, we use the fact that every complete consistent complete set has a saturated extension.

LEMMA [**Substitution**]. $\mathcal{M}[[\![x]\!] := [\![\pi]\!]] \models \phi$ *iff* $\mathcal{M} \models \phi[x := \pi]$. ■

LEMMA [**Succession**]. *If* $s \in \mathcal{S}$ *and* $\neg\Box\phi \in s$, *then there is a* $t \in \mathcal{S}$ *such that* $s \rightarrow_\Box t$ *and* $\neg\phi \in t$. ■

PROOF. Suppose that $s \in \mathcal{S}$ and $\neg\Box\phi \in s$. Let

$$\Phi = \{\psi \mid \Box\psi \in s\} \cup \{\neg\phi\}.$$

We show that Φ is (i) consistent and (ii) complete. Thus, a saturated extension t of Φ exists by part (ii) of the lemma on the *existence of saturated extensions*; furthermore, $t \in \mathcal{S}$ because $\approx_t = \approx_s$ by **RIG1-2**, and $s \rightarrow_\Box t$ by the definition of $\rightarrow_\Box$.

(i) Suppose that Φ is inconsistent, that is,

$$\vdash \psi_1 \rightarrow \cdots \rightarrow \psi_n \rightarrow \neg\phi \rightarrow \mathbf{false}$$

for some $\Box\psi_i \in s$, $1 \leq i \leq n$. By **K2** and repeated application of **K1**, $\Box(\neg\phi \rightarrow \mathbf{false}) \in s$, implying that $\Box\phi \in s$ (use **K**). Since also $\neg\Box\phi \in s$, it follows that s is inconsistent, a contradiction.

(ii) Assume that $\Phi \vdash \varphi[x := \pi]$ for all terms π, let x be such that $x \notin \phi$ and $\varphi' \in \varphi - x$; by **VAR2** it suffices to show that $\Phi \vdash \varphi'$. Similarly to part (i), by **K** it follows that $\Box(\neg\phi \rightarrow \varphi[x := \pi]) \in s$ for all π. Hence $\Box(\neg\phi \rightarrow \varphi') \in s$ by the completeness of s (if φ' equals $x.\varphi$, use $\mathbf{Q}^*$ and **K**), and $\neg\phi \rightarrow \varphi' \in \Phi$ by the definition of Φ. Since also $\neg\phi \in \Phi$, we conclude that $\Phi \vdash \varphi'$. ■

Thus we have shown that every consistent formula $\neg\phi$ is satisfiable, by the corresponding canonical model $\mathcal{M}(\neg\phi)$. By part (i) of the *consistency* lemma, it follows that $\vdash \phi$ for every valid formula ϕ. Recall that we have assumed an arbitrary normal logic Λ; thus K contains precisely all valid formulas.

We remark that it is an interesting open question under which conditions the general form of the quantifier introduction rule **Q2** can be replaced by the simpler rule $\mathbf{Q2}^*$. We conjecture that this is case for all normal logics Λ_Φ, such as K, whose set Φ of nonlogical axioms satisfies certain closure properties.[2]

[2]Note that the condition that Φ can be closed under **Q2** using only $\mathbf{Q2}^*$ is insufficient. Consider the normal logic with the single nonlogical axiom $p \rightarrow x = y$. Since this axiom contains no modal operators, all of its **Q2**-consequences are $\mathbf{Q2}^*$-consequences. Yet the formula $p \rightarrow \Box x.x = y$, which follows by **RIG1** and **Q2**, cannot be derived without the general rule **Q2**.

2.4 Syntactic and semantic extensions

In order to treat the real-time temporal logic TPTL as a half-order modal logic, we admit multiple modal operators and semantic restrictions on the corresponding accessibility relations.

If we restrict ourselves to interpretations $\mathcal{M}_C$ whose accessibility relation $\rightarrow_\Box$ satisfies a certain condition C, we have to add nonlogical axioms to our proof system so that all formulas that are true under all interpretations $\mathcal{M}_C$ can be derived.

For example, consider only interpretations with a reflexive accessibility relation. They clearly satisfy all formulas of the form $\Box\phi \rightarrow \phi$. In order to show that adding these formulas as nonlogical axioms is sufficient to characterize reflexive interpretations completely, simply observe that, with this axiom schema, the accessibility relation of the canonical model is reflexive.

The other parts of the following lemma are established similarly. (In fact, the proofs are identical to the corresponding arguments for propositional modal logic, and can, for example, be found in [Go87].)

LEMMA [**Accessibility conditions**]. *The following axiom schemata characterize the corresponding conditions on the accessibility relation $\rightarrow_\Box \subseteq \mathcal{S}^2$ completely.*

- *Reflexivity:* $\Box\phi \rightarrow \phi$.
- *Symmetry:* $\Box\phi \rightarrow \Box\Diamond\phi$.
- *Transitivity:* $\Box\phi \rightarrow \Box\Box\phi$.
- *Seriality and functionality* (for all $s \in \mathcal{S}$ there is a [unique] $t \in \mathcal{S}$ such that $s \rightarrow_\Box t$):

$$\Box\phi \rightarrow \Diamond\phi \text{ and } \Box\phi \leftrightarrow \Diamond\phi.$$

- *Weak connectivity* ($s \rightarrow_\Box t$ and $s \rightarrow_\Box r$ implies that $t \rightarrow_\Box r$ or $t = r$ or $r \rightarrow_\Box t$):

$$\Box(\phi_1 \wedge \Box\phi_1 \rightarrow \phi_2) \vee \Box(\phi_2 \wedge \Box\phi_2 \rightarrow \phi_1). \blacksquare$$

Half-order versions of *multimodal* logics are straightforwardly defined by admitting several modal operators $\Box_i$ in the syntax. An interpretation contains, accordingly, a separate accessibility relation $\rightarrow_{\Box_i} \subseteq \mathcal{S}^2$ for each operator $\Box_i$.

Multimodal normal logics are closed under the axiom schemata $\mathbf{K1}_i$, $\mathbf{RIG1}_i$, and $\mathbf{RIG2}_i$, as well as the inference rules $\mathbf{K2}_i$: one for each modal operator $\Box_i$. As for the inference rule **Q2**, consider every formula $\phi \rightarrow \Box_i\psi$, for any operator $\Box_i$, to be of the form $\phi \Rightarrow \psi$. For example, if $x \notin \phi_1, \phi_2$, then from

$$\phi_1 \rightarrow \Box_1(\phi_2 \rightarrow \Box_2\psi)$$

we may infer

$$\phi_1 \rightarrow \Box_1(\phi_2 \rightarrow \Box_2 x.\, \psi).$$

The completeness proof given above is easily generalized to show that multimodal K (the smallest multimodal normal logic) contains precisely all valid formulas.

2.5 Embedding classical logic

We have seen that half-order modal logic corresponds to a *fragment* of first-order *modal* logic, because the freeze quantifier can be expressed by conventional quantification in combination with a state variable. Alternatively, half-order modal logic can be viewed as a *generalization* of *classical* first-order logic. By showing how to embed classical first-order logic faithfully into a half-order normal logic, we prove the undecidability of the latter one.

Suppose that all predicates are rigid. Then we can read the combination "$\Box x.$" of a modal operator and the freeze quantifier as a universal quantifier with restricted scope: it ranges only over the values of adjacent states. Similarly, "$\Diamond x.$" can be viewed as a local existential quantifier. By pursuing this idea, we see that ordinary quantifiers are representable in a half-order modal logic with a universal accessibilty relation, that is, in whose models every state is accessible from every other state.

Let RIGID-S5 be the smallest normal logic containing the following nonlogical axioms:

SYMM $\quad \phi \rightarrow \Box\Diamond\phi$
TRANS $\quad \Box\phi \rightarrow \Box\Box\phi$
RIG3 $\quad p\bar{\pi} \rightarrow \Box p\bar{\pi}$ for all $p \in P$
EX $\quad \Diamond x.\, x = \pi.$

Note that in the presence of **EX**, which implies $\Diamond$ **true**, the schema

REFL $\quad \Box\phi \rightarrow \phi$

is derivable from **SYMM** and **TRANS**, while **SYMM** and **RIG4** imply

RIG4 $\quad \neg p\bar{\pi} \rightarrow \Box\neg p\bar{\pi}$ for all $p \in P$.

Together, the schemata **REFL**, **SYMM**, and **TRANS** characterize, as for the propositional modal logic S5, accessibility relations that are equivalence relations. The axiom schema **EX** assures that enough states are accessible (i.e., in the same equivalence class as the initial state). **RIG3** and **RIG4** assert that all predicate symbols are rigid.

Let ϕ be a classical formula over the first-order language (F, P), and ϕ^T the formula of half-order modal logic that results from replacing all quantifiers $\forall x.$ and $\exists x.$ by $\Box x.$ and $\Diamond x.$, respectively. The following theorem states that this translation preserves validity.

THEOREM [**Embedding of classical logic**]. *A classical first-order formula ϕ is valid iff ϕ^T is contained in* RIGID-S5. ∎

PROOF. First we show that

$$\phi^T \rightarrow \Box\phi^T \in \text{RIGID-S5} \qquad (\dagger)$$

for any classical first-order formula ϕ. We proceed by induction on the structure of ϕ, assuming that all negations in ϕ have been pushed inside in front of atomic formulas. If ϕ is an atomic formula or its negation, use one of **RIG1-4**. The propositional cases follow from the induction hypothesis by **K**. If ϕ is of the form $\forall x.\, \psi$, then

$$\Box x.\, \psi^T \rightarrow \Box\Box x.\, \psi^T$$

holds by **TRANS**. Finally, suppose that the outermost symbol of ϕ is an existential quantifier. In this case the inductive step is an instance of $\Diamond\phi \rightarrow \Box\Diamond\phi$, which follows from **SYMM** and **TRANS** by **K**.

Now assume that ϕ is provable by a complete Hilbert-style proof system for the first-order predicate calculus, say the one given in [En72]. Any classical deduction of ϕ (in the given proof system) can be transformed into a half-order modal derivation of ϕ^T, thus implying that $\phi^T \in$ RIGID-S5. The only interesting case is the derivation of the translation of the classical quantifier axiom $\psi[x := \pi] \rightarrow \exists x.\, \psi$ in RIGID-S5: from **EQ4**, by **K** infer

$$\Diamond x.\, x = \pi \rightarrow \Box\psi^T[x := \pi] \rightarrow \Diamond x.\, \psi^T,$$

which implies $\psi^T[x := \pi] \rightarrow \Diamond x.\, \psi^T$ by **EX** and (†).

The second direction of the theorem is shown semantically. Assume that $\phi^T \in$ RIGID-S5; that is, ϕ^T is true under all interpretations that satisfy the nonlogical axioms of RIGID-S5. We show that ϕ is true under an arbitrary classical first-order interpretation $\mathcal{I}$. Define the interpretation

$$\mathcal{M}_{\mathcal{I}} = (\mathcal{S}, \rightarrow_{\Box}, \mathcal{U}, |\,|, [\![x]\!]_{\mathcal{I}}, [\![f]\!]_{\mathcal{I}}, [\![p]\!], s_0)$$

for half-order modal logic such that

- both $\mathcal{S}$ and $\mathcal{U}$ are the universe of $\mathcal{I}$,
- $\rightarrow_{\Box} = \mathcal{S}^2$ is universal,
- $|\,|$ is the identity (i.e., $|s| = s$), and
- $[\![\bar{\pi}]\!] \in [\![p]\!](s)$ iff $[\![\bar{\pi}]\!] \in [\![p]\!]_{\mathcal{I}}$,

where $[\![x]\!]_{\mathcal{I}}$, $[\![f]\!]_{\mathcal{I}}$, and $[\![p]\!]_{\mathcal{I}}$ are the assignment functions of $\mathcal{I}$ for variables, function symbols, and predicate symbols, respectively. Any state of $\mathcal{M}_{\mathcal{I}}$ can be taken to be initial. It is not hard to see that $\mathcal{M}_{\mathcal{I}}$ satisfies all nonlogical axioms of RIGID-S5, and that $\mathcal{M}_{\mathcal{I}} \models \phi^T$ iff $\mathcal{I} \models \phi$. The theorem follows. ■

Since classical first-order logic is undecidable, so is RIGID-S5. This shows that axiomatizable half-order normal logics are not necessarily decidable.[3]

3 Timed Temporal Logic

In this section, we study half-order versions of the linear propositional temporal logic PTL (see, for example, [GPSS80]), which has been established as a working tool for the analysis of concurrent systems.

The semantics of half-order *temporal* logic is restricted to interpretations with a state structure that is isomorphic to the natural numbers NAT; these "temporal" interpretations are essentially infinite sequences of states. The syntax of half-order temporal logic contains two modal operators: the *next* operator $\bigcirc$, which is interpreted as "at the immediate successor state," and the *always* operator $\Box$ meaning "at all successor states" (i.e., "always in the future").

We show that, unlike in the propositional case, there cannot exist a complete half-order proof system for temporal structures in general.

Yet we introduced, in [AH89], a decidable half-order temporal logic, TPTL, for the specification and verification of real-time systems. The decidability of TPTL is due to a careful choice of both the set of values (monotonically increasing natural numbers) and the corresponding operations (the zero and successor functions, and the ordering predicate on NAT). Here we extend our proof system for half-order K to obtain a complete proof system for TPTL.

3.1 Syntax and semantics

The formulas of *timed temporal logic* (TPTL) are the formulas of half-order modal logic with the two modal operators $\bigcirc$ and $\Box$, where

[3] On the other hand, we suspect that every satisfiable formula ϕ of half-order logic is satisfiable under an interpretation that contains only a bounded finite number of states and a finite number of values (at most one for every state and one for every term in ϕ). It would follow that K itself is decidable in the half-order case.

- the set F of function symbols contains only the constant symbol 0 and the unary function symbol S (as in Successor), and
- the set P of predicate symbols contains only propositions P_0 (i.e., predicate symbols that take no arguments) and the binary predicate symbol $<$ (recall that equality is included in all half-order modal logics).[4]

Abbreviations such as $\leq$ and $+5$ are defined as usual.

TPTL is interpreted over timed state sequences, which are temporal structures whose values are monotonically increasing natural numbers. More precisely, a *timed state sequence* is an interpretation

$$\mathcal{M} = (\mathcal{S}, \rightarrow_{\bigcirc}, \rightarrow_{\Box}, \mathsf{NAT}, ||, [\![x]\!]_{x \in V}, [\![f]\!]_{f \in F}, [\![p]\!]_{p \in P}, \sigma_0)$$

for half-order modal logic such that

- $\rightarrow_{\bigcirc}$ imposes a linear order $\sigma_0 \rightarrow_{\bigcirc} \sigma_1 \rightarrow_{\bigcirc} \sigma_2 \rightarrow_{\bigcirc} \cdots$ on the set $\mathcal{S} = \{\sigma_i \mid i \geq 0\}$ of states,
- $\rightarrow_{\Box}$ is the reflexive transitive closure of $\rightarrow_{\bigcirc}$,
- $|\sigma_i| \leq |\sigma_{i+1}|$ for all $i \geq 0$,
- 0 denotes zero (i.e., $[\![0]\!] = 0$),
- S denotes the successor function on NAT (i.e., $[\![S]\!](n) = n + 1$), and
- $<$ rigidly denotes the ordering predicate on NAT (i.e., $[\![<]\!](s, m, n)$ iff $m < n$).

The correspondence with the standard definition of timed state sequences ([AH89], [AH90]) is obvious: any timed state sequence can be viewed as an infinite sequence of states $\sigma_i \subseteq P_0$, $i \geq 0$, together with an assignment function for the variables. Each of the states σ_i specifies the propositions that are true in that state (let $p \in \sigma_i$ iff $\sigma_i \in [\![p]\!]$), and has a value $|\sigma_i|$ associated with it. The values satisfy the *monotonicity* condition that, for all $i \geq 0$,

$$|\sigma_i| \leq |\sigma_{i+1}|.$$

This condition is motivated by the original design of TPTL as a real-time logic: the values that are associated with the states can be interpreted as time-stamps ([AH89]); think of state σ as representing the state of a system at time $|\sigma|$. From this point of view, the freeze quantifier "$x.$" binds the associated variable x to the "current" time.

Let us give a couple of examples for properties of timed state sequences that can be expressed in TPTL. The formula

$$x\,.\,\Box y\,.\,(y \leq x + 5 \;\rightarrow\; p)$$

asserts that p is true at all states within the next 5 time units, while

$$x\,.\,\Box y\,.\,(p \;\rightarrow\; y \leq x + 5)$$

specifies that p will not be true after 5 time units from now.

Note that, even though time is discrete, the notion of "next time" is entirely independent of "next state"; successive states may have the same or vastly different times associated with them, as long as the time does not decrease. This turns out to be crucial for the real-time analysis of asynchronous systems. If desired, the requirement that the time increases always by 1 between successive states, and thus acts as a state counter, can be expressed within TPTL, by the formula $\Box x.\ \bigcirc y\,.\,y = x + 1$.

[4] TPTL as originally defined in [AH89] differs syntactically in that all formulas have to be closed, and the freeze quantifiers are coupled with the temporal operators. We have already observed that this coupling does not restrict us in any essential way; recall the discussion of normal forms following the proof of **VAR3**.

3.2 Proof system

We extend the proof system for half-order K by axioms for timed state sequences.

The following three axiom schemata completely characterize temporal structures in propositional temporal logic ([GPSS80]):

LIN1 $\bigcirc\neg\phi \leftrightarrow \neg \bigcirc \phi$
LIN2 $\Box\phi \rightarrow \phi \wedge \bigcirc\Box\phi$
LIN3 $\phi \rightarrow \Box(\phi \rightarrow \bigcirc\phi) \rightarrow \Box\phi$.

Note that **LIN1** asserts the functionality of $\rightarrow_{\bigcirc}$, and that **LIN2** immediately implies **REFL**: $\Box\phi \rightarrow \phi$. **LIN3** gives an induction principle.

In addition, we need a set of axioms that allows us to derive all universal sentences of the decidable classical first-order theory of $(\mathsf{NAT}, 0, S, \leq)$. For instance, the following group of **NAT** axioms from [En72] has this completeness property, as can be shown by a quantifier elimination procedure (see [En72]):

NAT1 $x < Sy \leftrightarrow x \leq y$
NAT2 $x < y \vee x = y \vee y < x$
NAT3 $x < y \rightarrow y \not< x$
NAT4 $x < y \rightarrow y < z \rightarrow x < z$
NAT5 $x \not< 0$.

The axiom **MON** states that the time is monotonically increasing from state to state, and the axiom schema **PRO**[5] assures that all free variables are interpreted as finite natural numbers:

MON $x.\bigcirc y.x \leq y$
PRO $\Box x.\Diamond y.y > x \rightarrow \Diamond x.x > z$
RIG5 $\pi_1 < \pi_2 \rightarrow \bigcirc\pi_1 < \pi_2$.

RIG5 is sufficient to guarantee the rigidity of $<$ (see the upcoming lemma on sample theorems).

We show that TPTL is the smallest normal logic closed under the nonlogical axioms **LIN1-3**, **NAT1-5**, **MON**, **PRO**, and **RIG5**. It is straightforward to convince yourself that every axiom is true in all timed state sequences. Before proving completeness, we derive some additional formulas that will be useful later.

The following theorems can be proved purely propositionally (see, for example, [Go87] for the derivations):

TRANS $\Box\phi \rightarrow \Box\Box\phi$
LIN4 $\Box\phi \leftrightarrow \phi \wedge \bigcirc\Box\phi$
LIN5 $\Box(\Box\phi_1 \rightarrow \phi_2) \vee \Box(\Box\phi_2 \rightarrow \phi_1)$
LIN6 $\Box(\Box(\phi \rightarrow \Box\phi) \rightarrow \phi) \rightarrow \Diamond\Box\phi \rightarrow \Box\phi$.

LEMMA [**Sample theorems**]. *Any normal logic closed under the axioms given above contains the following formulas:*

MON′ $x.\Box y.x \leq y$
RIG6 $\pi_1 \not< \pi_2 \rightarrow \bigcirc\pi_1 \not< \pi_2$
RIG7 $\pi_1 < \pi_2 \rightarrow \Box\pi_1 < \pi_2$
RIG8 $\pi_1 \not< \pi_2 \rightarrow \Box\pi_1 \not< \pi_2$
TSS1 $\Box x.\bigcirc y.y = x \rightarrow x.\Box y.y = x$
TSS2 $x.\Box y.y = x \rightarrow (x.\Box\phi \leftrightarrow \Box x.\phi)$. ∎

[5] **PRO** has to be stated as an implication, because timed state sequences are, unlike in [AH89], not required to satisfy the *progress* condition that the time will never stagnate forever.

PROOF. **MON′** generalizes the monotonicity axiom **MON**. Its proof demonstrates the application of the induction schema **LIN3**. By applying $\mathbf{Q}^*$ to **LIN3**, it suffices to derive both $x.\, y.\, x \leq y$ and $x.\, \Box(y.\, x \leq y \rightarrow \bigcirc y.\, x \leq y)$. The base case follows from **QEQ** by $\mathbf{Q}^*$ (let us begin to suppress to mention applications of the equality axioms **EQ**).

To show the inductive step, by $\mathbf{Q2}^*$ and $\mathbf{K2}_\Box$ it suffices to derive $y.\, x \leq y \rightarrow \bigcirc y.\, x \leq y$. From **NAT**:

$$x \leq y \rightarrow y \leq z \rightarrow x \leq z.$$

By $\mathbf{Q}^*$ and $\mathbf{K}_\bigcirc$:

$$\bigcirc x \leq y \rightarrow \bigcirc z.\, y \leq z \rightarrow \bigcirc z.\, x \leq z.$$

By $\mathbf{RIG1\text{-}2}_\bigcirc$, **RIG5**, and $\mathbf{Q}^*$:

$$y.\, x \leq y \rightarrow y.\, \bigcirc z.\, y \leq z \rightarrow \bigcirc z.\, x \leq z.$$

The desired conclusion follows by **MON**, **REFL**, and **VAR2**.

RIG6 can be inferred from **NAT**:

$$\pi_1 \not< \pi_2 \leftrightarrow \pi_1 = \pi_2 \vee \pi_2 < \pi_1$$

by $\mathbf{RIG1}_\bigcirc$, **RIG5**, and $\mathbf{K}_\bigcirc$. **RIG7** (**RIG8**) follows from **LIN3** by **RIG5** (**RIG6**) and $\mathbf{K2}_\Box$. Note that the rigidity axioms $\mathbf{RIG1\text{-}2}_\Box$ are similarly derivable from $\mathbf{RIG1\text{-}2}_\bigcirc$, and thus can be omitted.

TSS1 is also derived by induction. By applying $\mathbf{Q}^*$ to **LIN3**, it suffices to derive the base case $x.\, y.\, y = x$, which holds by **QEQ**, and the inductive step

$$\Box x.\, \bigcirc y.\, y = x \rightarrow x.\, \Box(y.\, y = x \rightarrow \bigcirc y.\, y = x).$$

Using $\mathbf{Q}^*$, $\mathbf{K}_\Box$, and **VAR2**, it suffices to show that

$$y.\, \bigcirc z.\, z = y \rightarrow y.\, y = x \rightarrow \bigcirc z.\, z = x,$$

which follows from **EQ4** by $\mathbf{Q}^*$.

TSS2 follows from **EQ4** by $\mathbf{K}_\Box$ and $\mathbf{Q}^*$. ■

3.3 Updating time references

Let $\delta_{=k}$ ($\delta_{>K}$) be an abbreviation for the formula $x.\, \bigcirc y.\, y = x + k$ ($x.\, \bigcirc y.\, y > x + K$), which asserts that the time difference between the current state and its successor state is exactly k (greater than K, respectively).

These time-difference formulas will be used to update, in TPTL-formulas, references to the times of previous states. For example, the formula $x.\, \Box\psi$ holds in a state that contains $\delta_{=k}$ iff $x.\, \psi$ is true in that state and, intuitively speaking, "$x.\, \Box\psi[x := x - k]$" is true in its successor state. If k is greater than the number of successor symbols occurring in $x.\, \psi$, then the monotonicity of time can be exploited to simplify, to **true** or **false**, all timing constraints in "$x.\, \Box\psi[x := x - k]$" that refer to the current time x.

We write $x.\, \phi^k$ for the TPTL-formula that expresses the condition "$x.\, \phi[x := x - k]$"; $x.\, \phi^k$ results from $x.\, \phi$ by updating all references of ϕ to the current time x by the time difference k. For instance, if $x.\, \phi$ is $x.\, \Box y.\, (p \rightarrow y.\, y \leq x + 5)$, then $x.\, \phi^1$, $x.\, \phi^5$, and $x.\, \phi^6$ equal $x.\, \Box y.\, (p \rightarrow y.\, y \leq x + 4)$, $x.\, \Box y.\, (p \rightarrow y.\, y \leq x)$, and $x.\, \Box y.\, (p \rightarrow \mathbf{false})$, respectively.

Formally, we define $x.\phi^k$ inductively as follows: $x.\phi^0$ equals $x.\phi$, and $x.\phi^{k+1}$ results from $x.\phi^k$ by replacing, in ϕ^k, every free term $S^{c+1}x$ with S^cx, every subformula of the form $x = S^cy$, $S^cy = x$, or $S^cy < x$ with **false**, and every subformula $x < S^dy$ with **true**, provided that x is free and y is bound in ϕ^k.

The following lemma confirms that this definition updates time references correctly; it can be shown by structural induction.

LEMMA [**Time step**]. *$\mathcal{M}[x := |\sigma_0| - k] \models \phi$ iff $\mathcal{M} \models x.\phi^k$ for all timed state sequence $\mathcal{M}$ with $|\sigma_0| \geq k$.* ■

On the syntactic side, we have to make sure that our proof system is strong enough to derive the following validities on the updating of time references.

LEMMA [**More theorems**]. *Let K be the number of successor symbols in ϕ. Any normal logic closed under the axioms given above contains the following formulas (let x be different from y):*

UPD1	$y.\, y = x + k \rightarrow (\phi[y := x] \leftrightarrow y.\phi^k)$
UPD2	$y.\, y > x + K \rightarrow (\phi[y := x] \leftrightarrow y.\phi^K)$
UPD3	$\delta_{=k} \rightarrow (x.\, \bigcirc \phi \leftrightarrow \bigcirc x.\phi^k)$
UPD4	$\delta_{>K} \rightarrow (x.\, \bigcirc \phi \leftrightarrow \bigcirc x.\phi^K)$
CLOCK	$(\bigvee_{0 \leq k \leq K} \delta_{=k}) \vee \delta_{>K}$. ■

PROOF. **UPD1** and **UPD2** constitute the syntactic counterpart to the *time-step* lemma; they capture the essence of the definition of $x.\phi^k$. The proofs of **UPD1** and **UPD2** proceed by induction on the structure of ϕ.

Consider **UPD1**. By $\mathbf{Q}^*$ it suffices to derive

$$y = x + k \rightarrow (\phi[y := x] \leftrightarrow \phi^k).$$

The base cases follow from **NAT**. We present only the inductive step that introduces a new quantifier. By the condition of safe substitutivity, in

$$y = x + k \rightarrow ((z.\phi)[y := x] \leftrightarrow (z.\phi)^k)$$

z has to be different from both x and y; hence derive

$$y = x + k \rightarrow (z.\phi[y := x] \leftrightarrow z.\phi^k),$$

which follows from the induction hypothesis by $\mathbf{Q}^*$.

UPD3 follows from **UPD1** by $\mathbf{K}_\bigcirc$, $\mathbf{Q}^*$, and **VAR2** (rename y so that it does not occur in ϕ). **UPD4** follows similarly from **UPD2**.

CLOCK is, in fact, derivable for any constant $K \geq 0$ and *exclusive-or* connectives, implying that every state contains a unique time-difference formula. From **NAT**:

$$x \leq y \rightarrow (\bigvee_{0 \leq k \leq K} y = x + K) \vee y > y + K,$$

infer $x.\, \bigcirc y.x \leq y \rightarrow$ **CLOCK** by $\mathbf{Q}^*$, $\mathbf{K}_\bigcirc$, and **LIN1**. **CLOCK** follows by **MON** and **REFL**. ■

3.4 Completeness

We show that the proof system given above is complete for TPTL. The organization of the proof follows largely the completeness proof for the propositional case, PTL, as given in [Go87]. We proceed in two main steps.

Given a consistent TPTL-formula ϕ_0, the canonical-model construction does not directly provide a timed state sequence, because $\rightarrow_\Box$ is not the transitive closure of $\rightarrow_\bigcirc$. In order to have the induction axiom **LIN3** force $\rightarrow_\Box$ to be the transitive closure of $\rightarrow_\bigcirc$, we have to be able to characterize, by a TPTL-formula, all states that are reachable from the starting state by repeated traversal of $\rightarrow_\bigcirc$. This can be achieved by collapsing the states of the canonical model into a *finite* number of *finitely* representable states (i.e., finite, consistent sets of formulas). Our filtration process is derived from the tableau-decision procedure for TPTL ([AH89]).

While the structure $\mathcal{M}^F(\phi_0)$ that is obtained by filtration of the canonical model $\mathcal{M}(\phi_0)$ satisfies the desired transitive-closure property, it is still not a timed state sequence. The problem is that a state may have multiple successor states. Thus, in a second step, we unroll the states of $\mathcal{M}^F(\phi_0)$ into a timed state sequence $\mathcal{M}^T(\phi_0)$. If the unrolling is done carefully, in a way that preserves the truth of all eventualities, then the resulting "canonical" timed state sequence $\mathcal{M}^T(\phi_0)$ is, at last, a model for ϕ_0.

Let us assume that ϕ_0 does not any contain timing assertions that compare a variable with a natural number, and that ϕ_0 is closed (as are all TPTL-formulas as defined originally, in [AH89]). Absolute time references (such as $x = 5$ or $x + 1 > 3$) and free variables ("parameters") have to be treated with some care; we delay their discussion until later. We also assume that all bound variables in ϕ_0 are distinct; this can always be achieved by renaming.

Furthermore, let

$$\mathcal{M}(\phi_0) = (\mathcal{S}, \rightarrow_\bigcirc, \rightarrow_\Box, \mathcal{U}, ||, [\![x]\!]_{x \in V}, [\![f]\!]_{f \in F}, [\![p]\!]_{p \in P}, s_0)$$

be the canonical model for ϕ_0, as constructed in Section 2.3.

First, we remark that we may forget about the actual values (times) that are associated with the states in $\mathcal{M}(\phi_0)$, because the formulas $\delta_{=k}$ keep track of the time differences between adjacent states. In any timed state sequence, we can reconstruct the times from these time-difference formulas, modulo the initial time.

Let K be the number of successor symbols S occurring in ϕ_0. The key observation underlying the filtration is that we can restrict our attention to a *finite* number of time-difference formulas, namely $\delta_{=k}$ for all $0 \leq k \leq K$ and $\delta_{>K}$. This is because if the time difference between two states is larger than K, its actual value has no bearing on the truth of ϕ_0; thus every model of ϕ_0 can be compressed into a model all of whose time steps are at most $K + 1$ (simply reduce larger time steps to $K + 1$).

It follows that the truth of any TPTL-formula is determined by the truth of *finitely* many "subformulas." We define the closure $\mathit{Closure}(\phi_0)$ of ϕ_0 under subformulas as the smallest (i.e., finite) set containing $x.\,\phi_0$ (for some $x \notin \phi_0$) that is closed under the following operation Sub:

$$\begin{aligned}
&\mathit{Sub}(x.\,(\phi_1 \rightarrow \phi_2)) = \{x.\,\phi_1, x.\,\phi_2\}\\
&\mathit{Sub}(x.\,\bigcirc \phi) = \{x.\,\phi^i \mid 0 \leq i \leq K+1\}\\
&\mathit{Sub}(x.\,\Box\phi) = \{x.\,\phi, x.\,\bigcirc \Box\phi\}\\
&\mathit{Sub}(x.\,y.\,\phi) = \{x.\,\phi[y := x]\}.
\end{aligned}$$

Let the finite filtration set Γ contain all formulas in $\mathit{Closure}(\phi_0)$ as well as the time-difference formulas $\delta_{=i}$, $0 \leq i \leq K$, and $\delta_{>K}$. Note that the outermost symbol of every formula in Γ is a quantifier.

For $s, t \in \mathcal{S}$, let sFt iff $s \cap \Gamma = t \cap \Gamma$, and $s^F = \{t \mid sFt\}$. We overload the symbol $\in$ by also writing $\phi \in s^F$ iff $\phi \in s'$ for all $s'Fs$ (observe that F is an equivalence relation). Let

$$\mathcal{M}^F(\phi_0) = (\mathcal{S}^F, \rightarrow^F_{\bigcirc}, \rightarrow^F_{\Box}, [\![p]\!]^F_{p \in P'}, s^F_0)$$

be the state structure that results from the canonical model $\mathcal{M}(\phi_0)$ by ignoring the values and identifying all states that agree on Γ:

- $\mathcal{S}^F = \{t^F \mid t \in \mathcal{S} \text{ and } s_0 \rightarrow_{\Box} t\}$ (the reachable consistent subsets of Γ),
- $s^F \rightarrow^F_{\bigcirc} t^F$ iff $s' \rightarrow_{\bigcirc} t'$ for some $s'Fs$ and some $t'Ft$,
- $s^F \rightarrow^F_{\Box} t^F$ iff $s(\rightarrow^F_{\bigcirc})^n t$ for some $n \geq 0$ (the reflexive transitive closure of $\rightarrow^F_{\bigcirc}$),
- $[\![p]\!]^F(s^F)$ iff $p \in s'$ for all $s'Fs$.

Note that there are only finitely many states s^F, each of which can be uniquely identified by a characteristic formula $\phi(s) \in s$, which is a finite conjunction of formulas and negated formulas from Γ:

$$\bigwedge_{\phi \in \Gamma \cap s} \phi \wedge \bigwedge_{\phi \in \Gamma - s} \neg\phi.$$

Let $s^F \rightarrow^F_{\Diamond} t^F$ iff for all $s'Fs$, there is some $t'Ft$ such that $s' \rightarrow_{\Box} t'$. The following lemma is proved similarly to the propositional case ([Go87]).

LEMMA [**Filtration**].

(i) $\rightarrow^F_{\bigcirc}$ *is serial.*
(ii) $s \rightarrow_{\Box} t$ *implies* $s^F \rightarrow^F_{\Box} t^F$.
(iii) $\rightarrow^F_{\Box}$ *is reflexive, transitive, and connected (i.e.,* $s^F \rightarrow^F_{\Box} t^F$ *or* $t^F \rightarrow^F_{\Box} s^F$*).*
(iv) $\rightarrow^F_{\Diamond}$ *is reflexive, transitive, connected, and* $\rightarrow^F_{\Diamond} \subseteq \rightarrow^F_{\Box}$.

PROOF. *(i)* **LIN1** ensures the functionality of $\rightarrow_{\bigcirc}$ by the lemma on *accessibility conditions*; this implies the seriality of $\rightarrow^F_{\bigcirc}$. We remark that $\rightarrow^F_{\bigcirc}$ is, however, not functional.

(ii) Suppose that $s \rightarrow_{\Box} t$. Let ϕ be the finite disjunction of all characteristic formulas $\phi(r)$ with $s^F \rightarrow^F_{\Box} r^F$; clearly $\phi \in r$ iff $s^F \rightarrow^F_{\Box} r^F$. Use the induction schema **LIN3** to show that $\Box\phi \in s$, which implies that $\phi \in t$.

(iii) Reflexivity and transitivity hold by definition. The *accessibility-conditions* lemma implies that the relation $\rightarrow_{\Box}$ is reflexive because of **REFL**, transitive because of **TRANS**, and weakly connected because of **REFL** and **LIN5**; hence it is connected on all states reachable from s_0 by $\rightarrow_{\Box}$. The connectivity of $\rightarrow^F_{\Box}$ follows by part *(ii)*.

(iv) The reflexivity, transitivity, and connectivity of $\rightarrow^F_{\Diamond}$ follow from the corresponding properties of $\rightarrow_{\Box}$ (see part *(iii)*). Part *(ii)* implies that $\rightarrow^F_{\Diamond} \subseteq \rightarrow^F_{\Box}$. ∎

The *filtration* lemma implies that $\mathcal{M}^F(\phi_0)$ consists of a finite sequence of strongly connected $\rightarrow^F_{\Box}$-components, each one of which consists of a finite sequence of strongly connected $\rightarrow^F_{\Diamond}$-components. We will construct a temporal model for ϕ_0 by unrolling $\mathcal{M}^F(\phi_0)$ into an infinite sequence of states. This has to be in a way such that, whenever some state contains an eventuality $x.\Diamond\phi$, then it is satisfied in a state that is unrolled "later." The following lemma guarantees this property for the unrolling that maintains the order of strongly connected $\rightarrow^F_{\Diamond}$-components, and repeats all states in the final $\rightarrow^F_{\Box}$-component infinitely often.

LEMMA [**Unrolling**]. *Assume that* $\neg\phi$ *is equivalent to a formula in* Γ. *Let* $s \in \mathcal{S}$ *be such that* $t^F \not\rightarrow^F_{\Diamond} s^F$ *for some* t. *If* $\Box\phi \notin s$, *then either* $\phi \notin s$, *or* $\phi \notin t$ *for some* t *such that* $s^F \rightarrow^F_{\Box} t^F$ *and* $t^F \not\rightarrow^F_{\Diamond} s^F$. ∎

PROOF. Let $\neg\phi$ be equivalent to a formula in Γ; then, for all s,t with sFt, $\phi \in s$ iff $\phi \in t$ and, by $\mathbf{K}_\Box$, $\Box\phi \in s$ iff $\Box\phi \in t$. The proof proceeds as in the propositional case, using the *filtration* lemma and **LIN6** ([Go87]). ∎

The *unrolling* lemma allows us to unroll $\mathcal{M}^F(\phi_0)$ in the described fashion, into an infinite sequence of states σ_i, $i \geq 0$, such that

- $\sigma_0 = s_0^F$,
- $\sigma_i \rightarrow^F_{\bigcirc} \sigma_{i+1}$, and
- whenever $\sigma_i = s^F$, $\Diamond\phi \in s$, and ϕ is equivalent to some formula in Γ, then $\phi \in t$ for some t such that $\sigma_j = t^F$ for some $j \geq i$.

CLOCK implies that every state σ_i contains one of the time-difference formulas. We reassign, in accordance with the time-difference formulas, times to all the states, thus obtaining the *canonical timed state sequence*

$$\mathcal{M}^T(\phi_0) = (\mathcal{S}^T, \rightarrow^T_{\bigcirc}, \rightarrow^T_{\Box}, \mathsf{NAT}, |\,|^T, [\![x]\!]^T_{x\in V}, [\![p]\!]^T_{p\in P_0}, \sigma_0),$$

where

- $\mathcal{S}^T = \{\sigma_i \mid i \geq 0\}$,
- $\rightarrow^T_{\bigcirc} = \{(\sigma_i, \sigma_{i+1}) \mid i \geq 0\}$,
- $\rightarrow^T_{\Box}$ is the reflexive transitive closure of $\rightarrow_{\bigcirc}$,
- $|\sigma_{i+1}|^T = |\sigma_i|^T + k$ if $\delta_{=k} \in \sigma_i$, and
  $|\sigma_{i+1}|^T = |\sigma_i|^T + K + 1$ if $\delta_{>K} \in \sigma_i$,
- $[\![p]\!]^T(s^F)$ iff $[\![p]\!]^F(s^F)$.

Note that $|\sigma_0|^T$ and $[\![x]\!]^T$ are left arbitrary; they need to be specified only in case ϕ_0 contains any absolute time references or free variables, respectively.

The following main theorem asserts that we have indeed constructed a model of ϕ_0. The proof depends crucially on the *time-step* lemma as well as **UPD3** and **UPD4**, which ensure the consistency of all timing constraints.

THEOREM [**Canonical timed state sequence**]. *For all $i \geq 0$ and $\phi \in$ Closure(ϕ_0), $\phi \in \sigma_i$ iff $\mathcal{M}^T(\phi_0)[\sigma_0 := \sigma_i] \models \phi$.* ∎

PROOF. We apply induction on the structure of ϕ, for $\phi \in$ *Closure*(ϕ_0).

If ϕ is of the form $x.p$ for some proposition p, use **Q3**. The case that ϕ is of the form $x.m = n$, $x.m < n$, $x.x + m = x + n$, or $x.x + m < x + n$ follows from **NAT** and $\mathbf{Q}^*$. The propositional cases are established by $\mathbf{Q}^*$. If ϕ is of the form $x.y.\psi$, use the *substitution* lemma and **VAR3**.

Now suppose that ϕ is of the form $x.\bigcirc\psi$ and $|\sigma_{i+1}|^T = |\sigma_i|^T + k$; that is, $\delta_{=k} \in s$ if $k \leq K$, and $\delta_{>K} \in s$ otherwise. Furthermore, by the definition of $\rightarrow^F_{\bigcirc}$ there are s,t such that $\sigma_i = s^F$, $\sigma_{i+1} = t^F$, and $s \rightarrow_{\bigcirc} t$. Then, $\mathcal{M}^T(\phi_0)[\sigma_0 := \sigma_i] \models x.\bigcirc\psi$ iff $x.\psi^k \in t$ by the *time-step* lemma and the induction hypothesis. **LIN1** and the canonical-model construction imply that $x.\psi^k \in t$ iff $\bigcirc x.\psi^k \in s$; and $\bigcirc x.\psi^k \in s$ iff $x.\bigcirc\psi \in \sigma_i$ follows from **UPD3** or **UPD4**.

Finally, suppose that ϕ is of the form $x.\Box\psi$. Let $k_j = |\sigma_j|^T - |\sigma_i|^T$ for all $j \geq i$. In this case, $\mathcal{M}^T(\phi_0)[\sigma_0 := \sigma_i] \models x.\Box\psi$ iff $x.\psi^{k_j} \in \sigma_j$ for all $j \geq i$, by the *time-step* lemma and the induction hypothesis.

We show that $x.\Box\psi \in \sigma_i$ implies $x.\Box\psi^{k_j} \in \sigma_j$ for all $j \geq i$, by induction on j; then $x.\psi^{k_j} \in \sigma_j$ by **LIN4** and $\mathbf{Q}^*$. Assume that $x.\Box\psi^{k_j} \in \sigma_j$ and $\delta_{k'} \in \sigma_j$, and show that $x.\Box\psi^{k_j+k'} \in \sigma_{j+1}$. By the definition of $\rightarrow^F_{\bigcirc}$ there are s,t such that $\sigma_j = s^F$, $\sigma_{j+1} = t^F$, and $s \rightarrow_{\bigcirc} t$. From $x.\Box\psi^{k_j} \in s$,

by **LIN4** and $\mathbf{Q}^*$ $x.\bigcirc \Box\psi^{k_j} \in s$. Since $\delta_{k'} \in s$, by **UPD3** or **UPD4** $\bigcirc x.(\Box\psi^{k_j})^{k'} \in s$; that is, $\bigcirc x.\Box\psi^{k_j+k'} \in s$. Hence $x.\Box\psi^{k_j+k'} \in t$ by the canonical-model construction.

Conversely, assume that $x.\Diamond\psi \in \sigma_i$ and $x.\psi^{k_j} \notin \sigma_j$ for all $j \geq i$, and show a contradiction. First observe that, by induction on j, it follows that $x.\Diamond\psi^{k_j} \in \sigma_j$ for all $j \geq i$.

We distinguish two cases. If $k_j > K$ for some $j \geq i$, then $x.\Diamond\psi^K \in \sigma_j$ and $x.\psi^K \notin \sigma_{j'}$ for all $j' \geq j$. Let $\sigma_j = s^F$. Since $x \notin \Diamond\psi^K$, by $\mathbf{Q}^*$ $x.\Diamond\psi^K$ is equivalent to $\Diamond x.\psi^K$; hence, by the *unrolling* lemma there is some $j' \geq j$ with $\sigma_{j'} = t^F$ and $x.\psi^K \in t$, contradicting $x.\psi^K \notin \sigma_{j'}$.

On the other hand, suppose that there is some $j \geq i$ such that $k_{j'}$ is constant for all $j' \geq j$; then there is some such $j \geq i$ such that $\sigma_j = s^F$ and s^F is in the final $\rightarrow^F_\Box$-component of $\mathcal{M}^F(\phi_0)$. Therefore $\delta_{=0} \in t^F$ for all t^F with $s^F \rightarrow^F_\Box t^F$. By part (ii) of the *filtration* lemma, $\delta_{=0} \in t$ for all t such that $s \rightarrow_\Box t$; thus $\Box\delta_{=0} \in s$ by the *canonical-model* theorem. By **TSS1**, $x.\Box y.y = x \in s$, and by **TSS2**, $x.\Diamond\psi^{k_j}$ is equivalent to $\Diamond x.\psi^{k_j}$. Hence, by the *unrolling* lemma there is some $j' \geq j$ with $x.\psi^{k_j} \in \sigma_{j'}$, again a contradiction. ■

This finishes the completeness proof for TPTL. We conclude by indicating how absolute time references and free variables can be incorporated into our argument.

To handle absolute time references, we include in the filtration set Γ all formulas of the form $x.x = k$ for $0 \leq k \leq K$, as well as $x.x > K$. In the definition of the canonical timed state sequence $\mathcal{M}^T(\phi_0)$, let $|\sigma_0| = k$ if $x.x = k \in \sigma_0$, and $|\sigma_0| = K+1$ if $x.x > K \in \sigma_0$. The additional base cases in the proof of the main theorem can be shown by induction on the canonical state sequence.

If ϕ_0 contains free variables, then it is no longer the case that every model can be compressed into a model all of whose time steps are at most $K+1$. However, it is not hard to see that, if ϕ_0 contains N free variables, then ϕ_0 is satisfiable iff it is satisfiable by a timed state sequence all of whose time steps are at most $K' = (N+1)(K+1)$. This is because, in any interpretation, the difference between any two times that are either associated with a state or a free variable, can be reduced to $K+1$ without changing the truth of ϕ_0.

Thus the completeness proof goes through if we take the filtration set Γ large enough (replace K by K'), and include all formulas of the form $x = y + k$, $0 \leq k \leq K'$, and $x > y + K'$, for all free variables $x, y \in \phi_0$. The axiom **PRO**, which is not required for the derivation of any closed formula, ensures that the variable assignment function of the canonical timed state sequence can be defined properly.

We also remark that the addition of the temporal *until* operator $\mathcal{U}$ is straightforward. The syntax of TPTL can be extended to admit formulas of the form $\phi_1\mathcal{U}\phi_2$, which is interpreted as "ϕ_2 holds at some future state, and from now until then ϕ_1 is continuously true"; that is, for any timed state sequence $\mathcal{M}$,

> $\mathcal{M} \models \phi_1\mathcal{U}\phi_2$ iff $\mathcal{M}[\sigma_0 := \sigma_i] \models \phi_2$ for some $i \geq 0$, and $\mathcal{M}[s_0 := \sigma_j] \models \phi_1$ for all $0 \leq j < i$.

By adding the two axiom schemata

$$\begin{array}{ll} \textbf{UNTIL1} & \phi_1\mathcal{U}\phi_2 \rightarrow \Diamond\phi_2 \\ \textbf{UNTIL2} & \phi_1\mathcal{U}\phi_2 \leftrightarrow \phi_2 \vee (\phi_1 \wedge \bigcirc(\phi_1\mathcal{U}\phi_2)), \end{array}$$

which characterize the *until* operator completely in PTL ([GPSS80]), we obtain a complete proof system for TPTL with *until*.

3.5 Half-order temporal logic

We show that half-order temporal logic in general is, unlike TPTL, not (recursively) axiomatizable, and therefore highly undecidable.

From TPTL we obtain TPTL$^+$ by restraining time to provide a state counter (i.e., $\Box x.\bigcirc y.\, y = x + 1$), omitting $<$, and adding the binary function symbol $+$ that is interpreted as ordinary addition on NAT. In [AH89], it is proved that the validity problem for closed TPTL$^+$-formulas is Π^1_1-hard, which implies that there is no complete proof system for this logic, as well as for TPTL with addition. Here we show that if half-order temporal logic were axiomatizable, for any choice of function and predicate symbols, then so would be TPTL$^+$. It follows that temporal interpretations cannot be completely characterized in half-order modal logic.

Let ϕ^+ denote the conjunction of the following formulas:

COUNT1	$x.\, x = 0$
COUNT2	$\Box x.\bigcirc y.\, y = Sx$
SUCC1	$\Box x.\, Sx \neq 0$
SUCC2	$\Box x.\Box y.\,(Sx = Sy \rightarrow x = y)$
PLUS1	$\Box x.\, 0 + x = x$
PLUS2	$\Box x.\Box y.\, Sx + y = S(x + y)$
PLUS3	$\Box x.\Box y.\, x + y = y + x.$

Recall that a *temporal* interpretation for half-order modal logic is one in which the set of states together with the two accessibility relations $\rightarrow_{\bigcirc}$ and $\rightarrow_{\Box}$ is isomorphic to $(\mathsf{NAT}, +1, \leq)$ (the set of values as well as the interpretation of all function and predicate symbols is left arbitrary). The following theorem states that ϕ^+ completely characterizes addition in temporal interpretations.

THEOREM [TPTL$^+$]. *The closed* TPTL$^+$*-formula* ψ *is valid iff* $\phi^+ \rightarrow \psi$ *is true under all temporal interpretations of half-order modal logic.* ■

PROOF. It is not hard to see that the closed formula ϕ^+ is true precisely under all temporal interpretations in which the set of values contains NAT (modulo isomorphism), where $|\sigma_i| = i$ for all $i \geq 0$, and 0, S, and $+$ are, on NAT, interpreted as the zero, successor, and addition functions. The theorem follows. ■

Since TPTL$^+$ is Π^1_1-hard, so is the restriction of half-order modal logic to temporal interpretations. We have, in fact, shown that any extension of TPTL with a single uninterpreted binary function symbol is Π^1_1-hard, because this function symbol can be forced, by ϕ^+, to be interpreted as addition on NAT.

4 Applications

We demonstrate the use of the proof system for TPTL by deriving a practical proof rule for the verification of bounded-response properties of real-time systems. Then we discuss some possible time-independent applications of the freeze quantifier.

4.1 Real-time verification

In [AH89], we have shown how TPTL can be used to express real-time properties of reactive systems, and introduced tableau-based methods for the verification of TPTL-formulas over finite-state systems. The proof system given in this paper provides a syntactic verification tool that complements the semantic finite-state techniques.

Consider, for example, the following program consisting of two concurrent processes:

$$\{u = 0, v = 0\}$$

$$
\begin{array}{lll}
\alpha_1\colon\ \textbf{while}\ u = 0 & & \\
\quad \alpha_2\colon\ v := v + 1; & & \\
\alpha_3\colon\ \textbf{while}\ v \neq 0 & \| & \beta_1\colon\ u := 1; \\
\quad \alpha_4\colon\ v := v - 1; & & \beta_2. \\
\alpha_5. & &
\end{array}
$$

Let us assume that all tests are instantaneous and all assignments take, nondeterministically, at least 1 and at most n time units (for some constant n). This can be stated in TPTL by (an infinite number of) program axioms like

$$x.(at\text{-}\alpha_2 \wedge v_{=0} \rightarrow \Diamond y.(at\text{-}\alpha_1 \wedge v_{=1} \wedge x + 1 \leq y \leq x + n))$$

(we use propositions such as $at\text{-}\alpha_2$ and $v_{=0}$ to denote the locations of the program control and the value of the program variables).

The given program is interesting from the real-time perspective, because the faster, and thus more often, assignment α_2 is performed, the more time is consumed by a complete program execution. Using the complete proof system for TPTL, we can formally verify that the program terminates in at most $2n + n \cdot (n + 1)$ time units; that is, every timed state sequence that satisfies all program axioms is a model of the TPTL-formula

$$
\begin{array}{l}
(u_{=0} \wedge v_{=0} \wedge at\text{-}\alpha_1 \wedge at\text{-}\beta_1) \rightarrow \\
x.\Diamond y.(at\text{-}\alpha_5 \wedge at\text{-}\beta_2 \wedge y \leq x + n^2 + 3n).
\end{array}
$$

A typical step in such a deduction requires the chaining of bounded-response properties. Thus, the following inference rule turns out to be very useful:

$$
\begin{array}{lll}
\textbf{RESP} & \text{from} & \Box x.(\phi_1 \rightarrow \Diamond y.(\phi_2 \wedge y \leq x + m)) \\
 & \text{and} & \Box x.(\phi_2 \rightarrow \Diamond y.(\phi_3 \wedge y \leq x + n)) \\
 & \text{infer} & \Box x.(\phi_1 \rightarrow \Diamond y.(\phi_3 \wedge y \leq x + m + n))
\end{array}
$$

for all constants $m, n \geq 0$, provided that neither x nor y occurs free in any of ϕ_1, ϕ_2, and ϕ_3.

Since **RESP** is sound over all timed state sequences, by our completeness result it must be derivable within the given proof system for TPTL. However, this derivation, which we are going to sketch briefly, is extremely tedious; it vividly demonstrates the need of practical high-level inference rules for real-time verification.

So let us derive **RESP**. In fact, we show the stronger assertion that the two antecedents of **RESP** imply the consequent:

$$
\begin{array}{l}
\Box x.(\phi_1 \rightarrow \Diamond y.(\phi_2 \wedge y \leq x + m)) \rightarrow \\
\Box x.(\phi_2 \rightarrow \Diamond y.(\phi_3 \wedge y \leq x + n)) \rightarrow \\
\Box x.(\phi_1 \rightarrow \Diamond y.(\phi_3 \wedge y \leq x + m + n))
\end{array}
$$

is valid. By $\mathbf{K}_\Box$, **TRANS**, $\mathbf{Q}^*$, and **VAR2**, it suffices to derive

$$
\begin{array}{l}
x.\Diamond(\phi_2 \wedge y.y \leq x + m) \rightarrow \\
\Box(\phi_2 \rightarrow y.\Diamond(\phi_3 \wedge z.z \leq y + n)) \rightarrow \\
x.\Diamond(\phi_3 \wedge z.z \leq x + m + n),
\end{array}
$$

which can be rewritten (use $\mathbf{K}_\Box$) as

$$
\begin{array}{l}
\Box(y.\Box(z.z \leq y + n \rightarrow \neg\phi_3) \rightarrow \neg\phi_2) \rightarrow \\
x.\Box(z.z \leq x + m + n \rightarrow \neg\phi_3) \rightarrow \\
x.\Box(y.y \leq x + m \rightarrow \neg\phi_2).
\end{array}
$$

By $\mathbf{Q}^*$, $\mathbf{K}_\Box$, and **TRANS** show

$$\begin{array}{l}\Box(y.\,\Box(z.\,z \le y+n \;\rightarrow\; \neg\phi_3) \;\rightarrow\; \neg\phi_2) \;\rightarrow\\ \Box y.\,\Box(z.\,z \le x+m+n \;\rightarrow\; \neg\phi_3) \;\rightarrow\\ \Box(y.\,y \le x+m \;\rightarrow\; \neg\phi_2).\end{array}$$

Applying $\mathbf{K}_\Box$ and $\mathbf{Q}^*$ again, it suffices to derive

$$\begin{array}{l}\Box(z.\,z \le x+m+n \;\rightarrow\; \neg\phi_3) \;\rightarrow\\ y \le x+m \;\rightarrow\; \Box(z.\,z \le y+n \;\rightarrow\; \neg\phi_3).\end{array}$$

By $\mathbf{RIG1}_\Box$, **RIG7**, $\mathbf{K}^*$, and $\mathbf{Q}^*$, show

$$y \le x+m \;\rightarrow\; z \le y+n \;\rightarrow\; z \le x+m+n,$$

which follows from **NAT**.

In order to make real-time verification feasible for systems that lie outside the scope of decision procedures, the proof system for TPTL has to be not only extended by useful derivable rules such as **RESP**, but also complemented by

- a *program* part of proof rules that restrict the models under consideration to the execution sequences of a particular system (program), and
- a first-order *domain* part to reason about the underlying data domain.

Ideally, this will be done in accordance with a proof methodology for different classes of real-time properties (see, for example, [MP89] for the untimed case).

4.2 Final comments

While we have focused on the real-time application of half-order modal logic, there seem to be several other intriguing prospects for the use of the freeze quantifier.

Suppose that with every program state, we associate not a single time-stamp but an entire vector of the current values of all program variables, say $u_1, \ldots u_k$. The freeze quantifier "$x.$" binds this tuple to the variable x, and we have k functions *value-of-*u_i $(1 \le i \le k)$, one to access each component of x (i.e., the value of u_i). This allows us to assert program properties, such as u being increased by 1 in the next execution step:

$$x.\; \bigcirc\, y.\; \mathit{value\text{-}of\text{-}u}(y) = \mathit{value\text{-}of\text{-}u}(x) + 1$$

or, in half-order *dynamic* logic:

$$x.\, \langle u := u+1 \rangle\, y.\; \mathit{value\text{-}of\text{-}u}(y) = \mathit{value\text{-}of\text{-}u}(x) + 1.$$

These properties are ordinarily stated in a much richer, first-order, modal logic:

$$\forall x.\, (u = x \;\rightarrow\; \langle u := u+1 \rangle\, u = x+1).$$

It can be argued that most transition axioms and input-output conditions are more naturally written without universal quantifiers and auxiliary variables. The use of the freeze quantifier in program verification opens an entire new area of investigation; there may be similar trade-offs in expressiveness versus complexity as in the real-time case.

Another application of the freeze quantifier that comes to mind, concerns half-order logics of knowledge. Suppose that the set of values is a domain of agents, and that we have two modal operators, K for knowledge and $\Box$ to change the agent that is reasoning. Then "$x.$" denotes the personal pronoun "I," which cannot be modeled adequately in propositional epistemic logics. For instance, we can express properties such as "I know that everyone who I know knows that I am smarter than (s)he is":

$$x.\, K \Box y.\, K(x \textit{ smarter-than } y).$$

The formula $x.\, Ky.\, x = y$ captures the assertion that "I know who I am."

Acknowledgements. Many thanks to Rajeev Alur, Adam Grove, Zohar Manna, and Amir Pnueli for their continuous discussions and support.

References

[AH89] R. Alur, T.A. Henzinger, "A really temporal logic," 30th IEEE FOCS, 1989.

[AH90] R. Alur, T.A. Henzinger, "Real-time logics: complexity and expressiveness," 5th IEEE LICS, 1990.

[En72] H.B. Enderton, *A Mathematical Introduction to Logic*, Academic Press, 1972.

[Ga84] J.W. Garson, "Quantification in modal logic," *Handbook of Philosophical Logic, Vol. II* (D. Gabbay and F. Guenthner, eds.), Reidel, 1984.

[Go87] R. Goldblatt, *Logics of Time and Computation*, CSLI Lecture Notes No. 7, 1987.

[GPSS80] D. Gabbay, A. Pnueli, S. Shelah, J. Stavi, "On the temporal analysis of fairness," 7th ACM POPL, 1980.

[MP89] Z. Manna, A. Pnueli, "The anchored version of the temporal framework," *Linear Time, Branching Time, and Partial Order in Logics and Models for Concurrency* (J.W. deBakker, W.-P. de Roever, and G. Rozenberg, eds.), Springer *LNCS* **354**, 1989.

[Os90] J.S. Ostroff, *Temporal Logic for Real-time Systems*, Research Studies Press, 1989.

[PH88] A. Pnueli, E. Harel, "Applications of temporal logic to the specification of real-time systems," Formal Techniques in Real-time and Fault-tolerant Systems, Springer *LNCS* **331**, 1988.