

Linear-Time Temporal Answer Set Programming

FELICIDAD AGUADO and PEDRO CABALAR

University of Corunna, Spain
(e-mails: felicidad.aguado@udc.es, cabalar@udc.es)

MARTÍN DIÉGUEZ

Université d'Angers, France
(e-mail: martin.dieguezlodeiro@univ-angers.fr)

GILBERTO PÉREZ

University of Corunna, Spain
(e-mail: gilberto.perez@udc.es)

TORSTEN SCHAUB and ANNA SCHUHMAN

University of Potsdam, Germany
(e-mails: torsten@cs.uni-potsdam.de, anna.schuhmann@uni-potsdam.de)

CONCEPCIÓN VIDAL

University of Corunna, Spain
(e-mail: concepcion.vidal@udc.es)

submitted 14 September 2020; revised 8 November 2021; accepted 17 November 2021

Abstract

In this survey, we present an overview on (Modal) Temporal Logic Programming in view of its application to Knowledge Representation and Declarative Problem Solving. The syntax of this extension of logic programs is the result of combining usual rules with temporal modal operators, as in *Linear-time Temporal Logic* (LTL). In the paper, we focus on the main recent results of the non-monotonic formalism called *Temporal Equilibrium Logic* (TEL) that is defined for the full syntax of LTL but involves a model selection criterion based on *Equilibrium Logic*, a well known logical characterization of Answer Set Programming (ASP). As a result, we obtain a proper extension of the stable models semantics for the general case of temporal formulas in the syntax of LTL. We recall the basic definitions for TEL and its monotonic basis, the temporal logic of Here-and-There (THT), and study the differences between finite and infinite trace length. We also provide further useful results, such as the translation into other formalisms like Quantified Equilibrium Logic and Second-order LTL, and some techniques for computing temporal stable models based on automata constructions. In the remainder of the paper, we focus on practical aspects, defining a syntactic fragment called (*modal*) *temporal logic programs* closer to ASP, and explaining how this has been exploited in the construction of the solver **telingo**, a temporal extension of the well-known ASP solver **clingo** that uses its incremental solving capabilities.

KEYWORDS: answer set programming, linear-time temporal logic, equilibrium logic, non-monotonic reasoning

1 Introduction

The term *Temporal Logic Programming*¹ was introduced in the late 1980s to refer to an extension of logic programs that incorporates modal temporal operators, usually from Linear-time Temporal Logic or LTL (Kamp 1968; Pnueli 1977). These LTL extensions of logic programming include proposals by Moszkowski (1986), Fujita *et al.* (1986), Gabbay (1987b), Abadi and Manna (1989), Baudinet (1992) or Orgun and Wadge (1990). However, the initial interest gradually cooled down over time and LTL extensions have neither become a common feature nor had a substantial impact in Prolog. With the appearance of *Answer Set Programming* or ASP (Lifschitz 2002) and its use as a formalism for practical knowledge representation, the interest in the specification of dynamic systems was renewed. ASP has been commonly used for temporal reasoning and the representation of action theories, following the methodology proposed by Gelfond and Lifschitz (1993). This methodology represents transition systems by adding to all time-dependent predicates an extra integer parameter: the time point T ranging over the finite interval $0, 1, \dots, n-1$ where $n \in \mathbb{N}$ is some fixed length that can be iteratively increased. A major disadvantage of this representation is that it provides neither special language constructs nor specific inference methods for temporal reasoning, as available in LTL. As a result, it is infeasible to represent properties of reactive systems (with infinite runs or *traces*) such as *safety* (“Is some particular state reachable?”), *liveness* (“Does some condition happen infinitely often?”), or to conclude that a given planning problem has no solution at all. Even if we restrict ourselves to finite traces,² we cannot exploit in non-temporal ASP the succinctness of LTL expressions for characterizing plans of interest (say, sentences like “never execute A if $B \vee C$ has held since D ”, etc) or the interpretation of such expressions by their well-known translations into automata.

In this paper, we focus on the temporal extension of ASP with operators from LTL. Combinations of ASP with other temporal logics are outside the scope of this paper, although an overall comparison is given in Section 7. In particular, we give an overview of the main definitions and results for the (modal) Temporal Logic Programming formalism called *Temporal Equilibrium Logic* or TEL (Cabalar and Vega 2007). As suggested by its name, it combines LTL with the logical characterization of ASP based on *Equilibrium Logic* (Pearce 1997). Although TEL was introduced more than a decade ago and a first survey was already presented by Aguado *et al.* (2013), several important advances have taken place since then, providing significant breakthroughs. On the theoretical side, apart from new results on complexity and expressiveness (Bozzelli and Pearce 2015; Balbiani and Diéguez 2016; Aguado *et al.* 2017), the most important feature has arguably been the redefinition of TEL to cope not only with infinite traces, as in its original version, but also with finite traces (Cabalar *et al.* 2018), following similar steps to the same variation introduced for LTL by De Giacomo and Vardi (2013). On the practical side, finite traces are much better suited to the way in which ASP is used to solve planning problems and have paved the way for the introduction of temporal operators into the

¹ In this paper, we use this original meaning of the term, that is, the extension of *logic programming* with *temporal logic*. However, nowadays, temporal logic programming can be understood as the more general discipline of combining *logic programming* with *temporal* reasoning, possibly in other different ways. See related work in Section 7 for a more detailed explanation of this distinction.

² Note that, even for finite traces, LTL-satisfiability is still a PSPACE-complete problem, while ASP-satisfiability lies in Σ_2^F in the most general case.

ASP solver `clingo` (Gebser et al. 2019), giving birth to the first full-fledged temporal ASP solver called `telingo` (Cabalar et al. 2019). In the paper at hand, we give a revised definition of TEL that incorporates the new advances, present the most general version of the logic (which neither imposes nor forbids finiteness of traces) and then specify the two variants: TEL_ω for infinite traces, as first defined by Cabalar and Vega (2007), and TEL_f for finite ones, as recently introduced by Cabalar et al. (2018). We revisit previous results for TEL_ω under the more general umbrella of TEL and then compare their effects to the case of TEL_f , providing a homogeneous presentation of the main achievements in the topic.

The rest of the paper is organized as follows. In the next section, we introduce the semantic framework of TEL, whose basic definition is made in terms of a models selection criterion on top of the monotonic basis provided by the logic of *Temporal Here-and-There* (THT). We also explain the relation of TEL to standard ASP by proving that the former can be seen as a fragment of Quantified Equilibrium Logic with monadic predicates and a linear order relation. In Section 3, we study the monotonic basis, THT, in full detail, providing some properties and relations to other formalisms. Section 4 focuses on different techniques to compute temporal equilibrium models via an automata construction. This is done in two steps: first defining the temporal equilibrium models of a temporal theory in terms of a second-order formula expressed in the logic of *Quantified LTL* (QLTL), and then building the automata from this QLTL formula. In the next section, we study a normal form for temporal theories under TEL semantics: what we call (*modal*) *temporal logic programs*. This normal form is used as the basis for a pair of translations from temporal logic programs (for finite traces) into standard ASP programs. We also define a syntactic fragment, called *present-centered* rules, that facilitates the application of incremental reasoning and, eventually, has led to the construction of the tool `telingo`, an extension of `clingo` with temporal operators. Finally, Section 7 discusses related work and Section 8 concludes the paper.

2 Temporal Equilibrium Logic

The definition of TEL is done in two steps. First, we define a monotonic logic called (*Linear-time*) THT, a temporal extension of the intermediate logic of Here-and-There (Heyting 1930). In a second step, we select some models from THT that are said to be *in equilibrium*, obtaining in this way a non-monotonic entailment relation.

2.1 Monotonic basis: Temporal Here-and-There

The syntax of THT (and TEL) is the same as for LTL with past operators. In particular, in this paper, we use the following notation. Given a (countable, possibly infinite) set \mathcal{A} of propositional variables (called *alphabet*), *temporal formulas* φ are defined by the grammar:

$$\varphi ::= a \mid \perp \mid \varphi_1 \otimes \varphi_2 \mid \bullet\varphi \mid \varphi_1 \mathbf{S} \varphi_2 \mid \varphi_1 \mathbf{T} \varphi_2 \mid \circ\varphi \mid \varphi_1 \mathbf{U} \varphi_2 \mid \varphi_1 \mathbf{R} \varphi_2 \mid \varphi_1 \mathbf{W} \varphi_2$$

where $a \in \mathcal{A}$ is an atom and \otimes is any binary Boolean connective $\otimes \in \{\rightarrow, \wedge, \vee\}$. The last six cases correspond to the temporal connectives whose names are listed below:

<p><i>Past</i></p> <ul style="list-style-type: none"> ● for <i>previous</i> S for <i>since</i> T for <i>trigger</i> 	<ul style="list-style-type: none"> ○ for <i>next</i> U for <i>until</i> R for <i>release</i> W for <i>while</i>
--	---

the intended meaning of the previous temporal operators is the following: ● φ (resp. ○ φ) means that φ is true at the previous (resp. next) time point. $\varphi \cup \psi$ means that φ is true until ψ is true, while $\varphi \mathbf{S} \psi$ can be read as φ is true since ψ was true. For $\varphi \mathbf{R} \psi$ and $\varphi \mathbf{T} \psi$ the meaning is not as direct as for the previous operators. $\varphi \mathbf{R} \psi$ means that ψ is true until both φ and ψ become true simultaneously or ψ is true forever. $\varphi \mathbf{T} \psi$ means that ψ is true since both φ and ψ became true simultaneously or ψ has been true from the beginning.

We also define several common derived operators like the Boolean connectives $\top \stackrel{def}{=} \neg \perp$, $\neg \varphi \stackrel{def}{=} \varphi \rightarrow \perp$, $\varphi \leftrightarrow \psi \stackrel{def}{=} (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, and the following temporal operators:

$\blacksquare \varphi \stackrel{def}{=} \perp \mathbf{T} \varphi$ <i>always before</i>	$\square \varphi \stackrel{def}{=} \perp \mathbf{R} \varphi$ <i>always afterward</i>
$\blacklozenge \varphi \stackrel{def}{=} \top \mathbf{S} \varphi$ <i>eventually before</i>	$\diamond \varphi \stackrel{def}{=} \top \mathbf{U} \varphi$ <i>eventually afterward</i>
$\mathbf{I} \stackrel{def}{=} \neg \bullet \top$ <i>initial</i>	$\mathbf{F} \stackrel{def}{=} \neg \circ \top$ <i>final</i>
$\hat{\bullet} \varphi \stackrel{def}{=} \bullet \varphi \vee \mathbf{I}$ <i>weak previous</i>	$\hat{\circ} \varphi \stackrel{def}{=} \circ \varphi \vee \mathbf{F}$ <i>weak next</i>

A (*temporal*) *theory* is a (possibly infinite) set of temporal formulas. Note that we use solid operators to refer to the past, while future-time operators are denoted by outlined symbols.

Although THT and LTL share the same syntax, they have a different semantics, the former being a weaker logic than the latter. The semantics of LTL relies on the concept of a *trace*, a (possibly infinite) sequence of *states*, each of which is a set of atoms. For defining traces, we start by introducing some notation to deal with intervals of integer time points. Given $a \in \mathbb{N}$ and $b \in \mathbb{N} \cup \{\omega\}$, we let $[a..b]$ stand for the set $\{i \in \mathbb{N} \mid a \leq i \leq b\}$, $[a..b)$ for $\{i \in \mathbb{N} \mid a \leq i < b\}$ and $(a..b]$ for $\{i \in \mathbb{N} \mid a < i \leq b\}$. In LTL, a *trace* \mathbf{T} of length λ over alphabet \mathcal{A} is a sequence $\mathbf{T} = (T_i)_{i \in [0..\lambda)}$ of sets $T_i \subseteq \mathcal{A}$. We sometimes use the notation $|\mathbf{T}| \stackrel{def}{=} \lambda$ to stand for the length of the trace. We say that \mathbf{T} is *infinite* if $|\mathbf{T}| = \omega$ and *finite* if $|\mathbf{T}| \in \mathbb{N}$. To represent a given trace, we write a sequence of sets of atoms concatenated with ‘.’. For instance, the finite trace $\{a\} \cdot \emptyset \cdot \{a\} \cdot \emptyset$ has length 4 and makes a true at even time points and false at odd ones. For infinite traces, we sometimes use ω -regular expressions like, for instance, in the infinite trace $(\{a\} \cdot \emptyset)^\omega$ where all even positions make a true and all odd positions make it false.

At each state T_i in a trace, an atom a can only be true, viz. $a \in T_i$, or false, $a \notin T_i$. The logic THT weakens this truth assignment, following the same intuitions as the (non-temporal) logic of HT. In THT, an atom can have one of three truth-values in each state, namely, *false*, *assumed* (or true by default) or *proven* (or certainly true). Anything proved has to be assumed, but the opposite does not necessarily hold. Following this idea, a state i is represented as a pair of sets of atoms $\langle H_i, T_i \rangle$ with $H_i \subseteq T_i \subseteq \mathcal{A}$ where H_i (standing for “here”) contains the proven atoms, whereas T_i (standing for “there”) contains the assumed atoms. On the other hand, false atoms are just the ones not assumed, captured by $\mathcal{A} \setminus T_i$. Accordingly, a *Here-and-There trace* (for short HT-*trace*) of length λ over alphabet \mathcal{A} is a sequence of pairs $(\langle H_i, T_i \rangle)_{i \in [0..\lambda)}$ with $H_i \subseteq T_i$ for any $i \in [0..\lambda)$. For convenience, we usually represent the HT-trace as the pair $\langle \mathbf{H}, \mathbf{T} \rangle$ of traces $\mathbf{H} = (H_i)_{i \in [0..\lambda)}$ and $\mathbf{T} = (T_i)_{i \in [0..\lambda)}$. Given $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$, we also denote its length

as $|\mathbf{M}| \stackrel{def}{=} |\mathbf{H}| = |\mathbf{T}| = \lambda$. Note that the two traces \mathbf{H} , \mathbf{T} must satisfy a kind of order relation, since $H_i \subseteq T_i$ for each time point i . Formally, we define the ordering $\mathbf{H} \leq \mathbf{T}$ between two traces of the same length λ as $H_i \subseteq T_i$ for each $i \in [0..\lambda)$. Furthermore, we define $\mathbf{H} < \mathbf{T}$ as both $\mathbf{H} \leq \mathbf{T}$ and $\mathbf{H} \neq \mathbf{T}$. Thus, an HT-trace can also be defined as any pair $\langle \mathbf{H}, \mathbf{T} \rangle$ of traces such that $\mathbf{H} \leq \mathbf{T}$. The particular type of HT-traces satisfying $\mathbf{H} = \mathbf{T}$ are called *total*.

We proceed by generalizing the extension of HT with temporal operators, called THT (Aguado et al. 2013), to HT-traces of fixed length in order to integrate finite as well as infinite traces. Given any HT-trace $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$, we define the THT satisfaction of formulas as follows.

Definition 1 (THT-satisfaction; Aguado et al. 2013; Cabalar et al. 2018³)

An HT-trace $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$ of length λ over alphabet \mathcal{A} satisfies a temporal formula φ at time point $k \in [0..\lambda)$, written $\mathbf{M}, k \models \varphi$, if the following conditions hold:

1. $\mathbf{M}, k \models \top$ and $\mathbf{M}, k \not\models \perp$
2. $\mathbf{M}, k \models a$ if $a \in H_k$ for any atom $a \in \mathcal{A}$
3. $\mathbf{M}, k \models \varphi \wedge \psi$ iff $\mathbf{M}, k \models \varphi$ and $\mathbf{M}, k \models \psi$
4. $\mathbf{M}, k \models \varphi \vee \psi$ iff $\mathbf{M}, k \models \varphi$ or $\mathbf{M}, k \models \psi$
5. $\mathbf{M}, k \models \varphi \rightarrow \psi$ iff $\langle \mathbf{H}', \mathbf{T} \rangle, k \not\models \varphi$ or $\langle \mathbf{H}', \mathbf{T} \rangle, k \models \psi$, for all $\mathbf{H}' \in \{\mathbf{H}, \mathbf{T}\}$
6. $\mathbf{M}, k \models \bullet\varphi$ iff $k > 0$ and $\mathbf{M}, k-1 \models \varphi$
7. $\mathbf{M}, k \models \varphi \mathbf{S} \psi$ iff for some $j \in [0..k]$, we have $\mathbf{M}, j \models \psi$ and $\mathbf{M}, i \models \varphi$ for all $i \in (j..k]$
8. $\mathbf{M}, k \models \varphi \mathbf{T} \psi$ iff for all $j \in [0..k]$, we have $\mathbf{M}, j \models \psi$ or $\mathbf{M}, i \models \varphi$ for some $i \in (j..k]$
9. $\mathbf{M}, k \models \circ\varphi$ iff $k + 1 < \lambda$ and $\mathbf{M}, k+1 \models \varphi$
10. $\mathbf{M}, k \models \varphi \cup \psi$ iff for some $j \in [k..\lambda)$, we have $\mathbf{M}, j \models \psi$ and $\mathbf{M}, i \models \varphi$ for all $i \in [k..j)$
11. $\mathbf{M}, k \models \varphi \mathbb{R} \psi$ iff for all $j \in [k..\lambda)$, we have $\mathbf{M}, j \models \psi$ or $\mathbf{M}, i \models \varphi$ for some $i \in [k..j)$
12. $\mathbf{M}, k \models \varphi \mathbb{W} \psi$ iff for all $j \in [k..\lambda)$, we have $\langle \mathbf{H}', \mathbf{T} \rangle, j \models \varphi$ or $\langle \mathbf{H}', \mathbf{T} \rangle, i \not\models \psi$ for some $i \in [k..j)$ and for all $\mathbf{H}' \in \{\mathbf{H}, \mathbf{T}\}$

⊗

In general, these conditions inherit the interpretation of connectives from LTL (with past operators) with just a few differences. A first minor variation is that we allow traces of arbitrary length λ , including both infinite ($\lambda = \omega$) and finite ($\lambda \in \mathbb{N}$) traces. The most significant difference, however, has to do with the treatment of implication, which is inherited from the intermediate logic of HT. It requires that the implication is satisfied in “both dimensions” \mathbf{H} (here) and \mathbf{T} (there) of the trace, using $\langle \mathbf{H}, \mathbf{T} \rangle$ (as in the other connectives) but also $\langle \mathbf{T}, \mathbf{T} \rangle$. Finally, one last difference with respect to LTL is the new connective $\varphi \mathbb{W} \psi$ which is also a kind of temporally iterated HT implication. Its intuitive reading⁴ is “keep doing φ while condition ψ holds.” In LTL, $\varphi \mathbb{W} \psi$ would just amount to $\neg\psi \mathbb{R} \varphi$, but under HT semantics both formulas have a different meaning, as the latter may provide evidence for φ even though the condition ψ does not hold.

³ The while operator \mathbb{W} was introduced by Aguado et al. (2020).

⁴ A dual, past operator could also be easily defined, but it would not have a natural, intuitive reading and its purpose would be unclear.

An HT-trace \mathbf{M} is a *model* of a temporal theory Γ if $\mathbf{M}, 0 \models \varphi$ for all $\varphi \in \Gamma$. We write $\text{THT}(\Gamma, \lambda)$ to stand for the set of THT-models of length λ of a theory Γ , and define $\text{THT}(\Gamma) \stackrel{\text{def}}{=} \text{THT}(\Gamma, \omega) \cup \bigcup_{\lambda \in \mathbb{N}} \text{THT}(\Gamma, \lambda)$. That is, $\text{THT}(\Gamma)$ is the whole set of models of Γ of any length. For $\Gamma = \{\varphi\}$, we just write $\text{THT}(\varphi, \lambda)$ and $\text{THT}(\varphi)$. We can analogously define $\text{LTL}(\Gamma, \lambda)$, that is, the set of traces of length λ that satisfy theory Γ , and $\text{LTL}(\Gamma)$, that is, the LTL-models of Γ any length. We omit specifying LTL satisfaction since it coincides with THT when HT-traces are total.

Proposition 1 (Aguado et al. 2013; Cabalar et al. 2018)

Let \mathbf{T} be a trace of length λ , φ a temporal formula, and $k \in [0.. \lambda)$ a time point.

Then, $\mathbf{T}, k \models \varphi$ in LTL iff $\langle \mathbf{T}, \mathbf{T} \rangle, k \models \varphi$. \(\square\)

In fact, total models can be forced by adding the following set of *excluded middle* axioms:

$$\square(a \vee \neg a) \quad \text{for each atom } a \in \mathcal{A} \text{ in the alphabet.} \quad (\text{EM})$$

Proposition 2 (Aguado et al. 2013; Cabalar et al. 2018)

Let $\langle \mathbf{H}, \mathbf{T} \rangle$ be an HT-trace and (EM) the theory containing all excluded middle axioms for every atom $a \in \mathcal{A}$. Then, $\langle \mathbf{H}, \mathbf{T} \rangle$ is a model of (EM) iff $\mathbf{H} = \mathbf{T}$. \(\square\)

Satisfaction of derived operators can be easily deduced, as shown next.

Proposition 3 (Aguado et al. 2013; Cabalar et al. 2018)

Let $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$ be an HT-trace of length λ over \mathcal{A} . Given the respective definitions of derived operators, we get the following satisfaction conditions:

12. $\mathbf{M}, k \models \mathbf{I}$ iff $k = 0$
13. $\mathbf{M}, k \models \widehat{\bullet}\varphi$ iff $k = 0$ or $\mathbf{M}, k-1 \models \varphi$
14. $\mathbf{M}, k \models \blacklozenge\varphi$ iff $\mathbf{M}, i \models \varphi$ for some $i \in [0..k]$
15. $\mathbf{M}, k \models \blacksquare\varphi$ iff $\mathbf{M}, i \models \varphi$ for all $i \in [0..k]$
16. $\mathbf{M}, k \models \mathbb{F}$ iff $k + 1 = \lambda$
17. $\mathbf{M}, k \models \widehat{\circ}\varphi$ iff $k + 1 = \lambda$ or $\mathbf{M}, k+1 \models \varphi$
18. $\mathbf{M}, k \models \diamond\varphi$ iff $\mathbf{M}, i \models \varphi$ for some $i \in [k.. \lambda)$
19. $\mathbf{M}, k \models \square\varphi$ iff $\mathbf{M}, i \models \varphi$ for all $i \in [k.. \lambda)$

\(\square\)

A formula φ is a *tautology* (or is *valid*), written $\models \varphi$, iff $\mathbf{M}, k \models \varphi$ for any HT-trace \mathbf{M} and any $k \in [0.. \lambda)$. We call the logic induced by the set of all tautologies *Temporal logic of Here-and-There* (THT for short).

Several types of equivalence can be defined in THT. In this sense, it is important to observe that being equivalent is something generally stronger than simply having the same set of models. Two formulas φ, ψ are (globally) *equivalent*, written $\varphi \equiv \psi$, iff $\models \varphi \leftrightarrow \psi$, that is, $\mathbf{M}, k \models \varphi \leftrightarrow \psi$ for any HT-trace \mathbf{M} of length λ and any $k \in [0.. \lambda)$. THT-equivalence satisfies the rule of substitution of equivalents:

Proposition 4 (Substitution of equivalents)

Let $\gamma[\varphi]$ be a formula with some occurrence of subformula φ and $\gamma[\psi]$ denote the replacement of that occurrence by ψ in $\gamma[\varphi]$. If $\varphi \equiv \psi$ then $\gamma[\varphi] \equiv \gamma[\psi]$.

On the other hand, we say that φ, ψ are just *initially equivalent*, written $\varphi \equiv_0 \psi$, if they have the same models $\text{THT}(\varphi) = \text{THT}(\psi)$, that is, $\mathbf{M}, 0 \models \varphi$ iff $\mathbf{M}, 0 \models \psi$, for any HT-trace \mathbf{M} . Obviously, $\varphi \equiv \psi$ implies $\varphi \equiv_0 \psi$ but not vice versa. For example, note that $\bullet a \equiv_0 \perp$, since $\bullet a$ is always false at the initial situation, whereas in the general case $\bullet a \not\equiv \perp$ or, otherwise, we could always replace $\bullet a$ by \perp in any context.

One important remark is that the finiteness of a trace $\langle \mathbf{H}, \mathbf{T} \rangle$ only affects the satisfaction of formulas dealing with future-time operators. In particular, if $\langle \mathbf{H}, \mathbf{T} \rangle$ has some finite length $\lambda = n$, then, in the semantics for \mathbb{U} , \mathbb{R} and \mathbb{W} in Definition 1, index j ranges over the finite interval $\{k, \dots, n-1\}$. Besides, if $\lambda = n$, the satisfaction of \circ forces $k < n$, which implies that there does exist a next state $k + 1$. As a result, the formula $\circ\top$ is not always satisfied, since it is false whenever $k = n = \lambda$.

Operators \mathbf{I} and \mathbb{F} exclusively depend on the value of time point k , so that the valuation of atoms in $\langle \mathbf{H}, \mathbf{T} \rangle$ is irrelevant to them. As a result, they behave “classically” and satisfy the law of the excluded middle, that is, $\mathbf{I}\vee\neg\mathbf{I}$ and $\mathbb{F}\vee\neg\mathbb{F}$ are THT tautologies. Besides, operator \mathbb{F} can only be true in finite traces. This implies that the inclusion of axiom $\diamond\mathbb{F}$ in any theory forces its models to be finite traces, while including its negation $\neg\diamond\mathbb{F}$ causes the opposite effect, that is, all models of the theory are infinite traces.

Several logics stronger than THT can be obtained by the addition of axioms (or the corresponding restriction on sets of traces). For instance, THT_ω is defined as THT plus $\neg\diamond\mathbb{F}$, that is, THT where we exclusively consider infinite HT-traces.⁵ THT_f , the finite-trace version, corresponds to THT plus $\diamond\mathbb{F}$. Linear Temporal Logic for possibly infinite traces, LTL, can be obtained as THT plus (EM), that is, THT with total HT-traces, LTL_ω is captured by THT_ω plus (EM), that is, infinite and total HT-traces, and finally LTL_f can be obtained as THT_f plus (EM), that is, LTL on finite traces (De Giacomo and Vardi 2013).

We study more properties and results about THT later on, but we proceed next to define its non-monotonic extension, TEL.

2.2 Non-monotonic extension: Temporal Equilibrium logic

Given a set of THT-models, we define the ones in equilibrium as follows.

Definition 2 (Temporal Equilibrium/Stable Model)

Let \mathfrak{S} be some set of HT-traces.

A total HT-trace $\langle \mathbf{T}, \mathbf{T} \rangle \in \mathfrak{S}$ is a *temporal equilibrium model* of \mathfrak{S} iff there is no other $\mathbf{H} < \mathbf{T}$ such that $\langle \mathbf{H}, \mathbf{T} \rangle \in \mathfrak{S}$.

The trace \mathbf{T} is called a *temporal stable model* (TS-model) of \mathfrak{S} . □

We further talk about temporal equilibrium or temporal stable models of a theory Γ when $\mathfrak{S} = \text{THT}(\Gamma)$, respectively. Moreover, we write $\text{TEL}(\Gamma, \lambda)$ and $\text{TEL}(\Gamma)$ to stand for the temporal equilibrium models of $\text{THT}(\Gamma, \lambda)$ and $\text{THT}(\Gamma)$ respectively. The corresponding sets of TS-models are denoted as $\text{TSM}(\Gamma, \lambda)$ and $\text{TSM}(\Gamma)$ respectively. One interesting observation is that, since temporal equilibrium models are total models $\langle \mathbf{T}, \mathbf{T} \rangle$, due to Proposition 1, we obtain $\text{TSM}(\Gamma, \lambda) \subseteq \text{LTL}(\Gamma, \lambda)$ that is, temporal stable models are a subset of LTL-models.

⁵ This corresponds to the (stronger) version of THT considered previously by Aguado et al. (2013).

Since the ordering relation among traces is only defined for a fixed λ , the following can be easily observed:

Proposition 5 (Cabalar et al. 2018)

The set of temporal equilibrium models of Γ can be partitioned by the trace length λ , that is, $\bigcup_{\lambda=0}^{\omega} \text{TEL}(\Gamma, \lambda) = \text{TEL}(\Gamma)$. \square

TEL is the (non-monotonic) logic induced by temporal equilibrium models. We can also define the variants TEL_{ω} and TEL_f by applying the corresponding restriction to infinite and finite traces, respectively.

As an example of non-monotonicity, consider the formula

$$\square(\bullet\text{loaded} \wedge \neg\text{unloaded} \rightarrow \text{loaded}), \tag{1}$$

that corresponds to the inertia for *loaded*, together with the fact *loaded*, describing the initial state for that fluent. Without entering into too much detail, this formula behaves as the logic program with the rules:

`loaded(0).`

`loaded(T) :- loaded(T-1), not unloaded(T).`

for any time point $T > 0$. As expected, for some fixed λ , we get a unique temporal stable model of the form $\{\text{loaded}\}^{\lambda}$. This entails that *loaded* is always true, viz. $\square\text{loaded}$, as there is no reason for *unloaded* to become true. Note that in the most general case of TEL, we actually get one stable model per each possible λ , including $\lambda = \omega$. Now, consider formula (1) along with $\text{loaded} \wedge \circ\circ\text{unloaded}$ which amounts to adding the fact `unloaded(2)`. As expected, for each λ , the only temporal stable model now is $\mathbf{T} = \{\text{loaded}\} \cdot \{\text{loaded}\} \cdot \{\text{unloaded}\} \cdot \emptyset^{\alpha}$ where α can be $*$ or ω . Note that by making $\circ\circ\text{unloaded}$ true, we are also forcing $|\mathbf{T}| \geq 3$, that is, there are no temporal stable models (nor even THT-models) of length smaller than three. Thus, by adding the new information $\circ\circ\text{unloaded}$ some conclusions that could be derived before, such as $\square\text{loaded}$, are not derivable any more.

As an example emphasizing the behavior of finite traces, take the formula

$$\square(\neg a \rightarrow \circ a), \tag{2}$$

which can be seen as a program rule “`a(T+1) :- not a(T)`” for any natural number T . As expected, temporal stable models make *a* false in even states and true in odd ones. However, we cannot take finite traces making *a* false at the final state $\lambda - 1$, since the rule would force $\circ a$ and this implies the existence of a successor state. As a result, the temporal stable models of this formula have the form $(\emptyset \cdot \{a\})^+$ for finite traces in TEL_f , or the infinite trace $(\emptyset \cdot \{a\})^{\omega}$ in TEL_{ω} .

Another interesting example is the temporal formula

$$\square(\neg\circ a \rightarrow a) \wedge \square(\circ a \rightarrow a).$$

The corresponding rules “`a(T) :- not a(T+1)`” and “`a(T) :- a(T+1)`” have no stable model (Fages 1994) when grounded for all natural numbers T . This is because there is no way to build a finite proof for any `a(T)`, as it depends on infinitely many next states to be evaluated. The same happens in TEL_{ω} , that is, we get no infinite temporal stable model. However in TEL_f , we can use the fact that $\circ a$ is always false in the last state. Then, $\square(\neg\circ a \rightarrow a)$ supports *a* in that state and therewith $\square(\circ a \rightarrow a)$ inductively supports *a* everywhere.

As an example of a temporal expression not so close to logic programming, consider the formula $\Box\Diamond a$, which is normally used in LTL_ω to assert that a occurs infinitely often. As discussed by De Giacomo and Vardi (2013), if we assume finite traces, then the formula collapses to $\Box(\mathbb{F} \rightarrow a)$ in LTL_f , that is, a is true at the final state (and either true or false everywhere else). The same behavior is obtained in THT_ω and THT_f , respectively. However, if we move to TEL, a truth minimization is additionally required. As a result, in TEL_f , we obtain a unique temporal stable model for each fixed $\lambda \in \mathbb{N}$, in which a is true at the last state, and false everywhere else. Unlike this, TEL_ω yields no temporal stable model at all. This is because for any \mathbf{T} with an infinite number of a 's we can always take some \mathbf{H} from which we remove a at some state, and still have an infinite number of a 's in \mathbf{H} . Thus, for any total THT_ω -model $\langle \mathbf{T}, \mathbf{T} \rangle$ of $\Box\Diamond a$ there always exists some model $\langle \mathbf{H}, \mathbf{T} \rangle$ with strictly smaller $\mathbf{H} < \mathbf{T}$. Note that we can still specify infinite traces with an infinite number of occurrences of a , but at the price of removing the truth minimization for that atom. This can be done, for instance, by adding the excluded middle axiom (EM) for atom a . In this way, infinite traces satisfying $\Box\Diamond a \wedge \Box(a \vee \neg a)$ are those that contain an infinite number of a 's. In fact, if we add the excluded middle axiom for all atoms, TEL collapses into LTL, as stated below.

Proposition 6 (Aguado et al. 2013; Cabalar et al. 2018)

Let Γ be a temporal theory over \mathcal{A} and (EM) be the set of all excluded middle axioms for all atoms in \mathcal{A} .

Then, $TSM(\Gamma \cup (EM)) = LTL(\Gamma)$. ⊠

2.3 Relation to ASP

As we have seen in the above examples, there seems to be a connection between temporal formulas over propositional atoms like a and logic programs with atoms for (monadic) predicates, viz. $\mathbf{a}(\mathbf{T})$, with an integer argument \mathbf{T} representing a time point. This connection is strongly related to Kamp's well-known theorem (Kamp 1968) that allows for translating LTL into Monadic First-Order Logic with a linear order $<$ relation, $MFO(<)$ for short. In this section, we show how Kamp's translation is also applicable to TEL, so that the latter can also be reduced to (Monadic) Quantified Equilibrium Logic, which essentially covers ASP for predicates with one argument. This connection reinforces the adequacy of TEL as a suitable temporal extension of ASP.

We begin by revisiting the definition of *Quantified Equilibrium Logic* or QEL Pearce and Valverde (2008). This logic allows for first-order logic programs in ASP to be partially simplified before grounding. Moreover, its monotonic basis, *Quantified Here-and-There* (QHT), can be used to check the property of strong equivalence, analogous to HT in the propositional case. The definition of QHT is based on a first-order language denoted by $\langle C, F, P \rangle$, where C , F and P are three disjoint sets representing constants, functions and predicates, respectively. Given a set of constants D , we define:

- $\mathcal{T}(D, F)$ as the set of all ground terms that can be built with functions in F and constants in D , and
- $\mathcal{A}(D, P)$ as the set of all ground atomic sentences that can be formed with predicates in P and constants in D .

In its most general version, QHT allows for different domains in the here and there worlds and the interpretation of equality can also be varied in each world. In this paper, though, we use the most common option when applied to ASP, that is, QHT with so-called *static domains and decidable equality*, dealing with a common universe and a fixed interpretation of equality. In this setting, a QHT-interpretation is a tuple $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ such that:

- D is a (possibly infinite) non-empty set of constant names identifying each element in the universe. For simplicity, we use the same name for the constant and the universe element.
- $\sigma : \mathcal{T}(C \cup D, F) \rightarrow D$ is a mapping from ground terms to elements of D satisfying $\sigma(d) = d$ for all $d \in D$ and structural recursion $\sigma(f(t_1, \dots, t_n)) = f(\sigma(t_1), \dots, \sigma(t_n))$.
- H and T are sets of atomic sentences satisfying $H \subseteq T \subseteq \mathcal{A}(D, P)$.

Given two QHT-interpretations, $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ and $\mathcal{M}' = \langle (D', \sigma'), H', T' \rangle$, we say that $\mathcal{M} \leq \mathcal{M}'$ iff $D = D'$, $\sigma = \sigma'$, $T = T'$ and $H \subseteq H'$. If, additionally, $H \subset H'$ we say that the relation is strict and denote it by $\mathcal{M} < \mathcal{M}'$.

Definition 3 (QHT-satisfaction; Pearce and Valverde 2008)

A QHT-interpretation $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ satisfies a first-order formula α , written $\mathcal{M} \models \alpha$, if the following conditions hold:

- $\mathcal{M} \models \top$ and $\mathcal{M} \not\models \perp$
- $\mathcal{M} \models p(\tau_1, \dots, \tau_n)$ iff $p(\sigma(\tau_1), \dots, \sigma(\tau_n)) \in H$
- $\mathcal{M} \models \tau = \tau'$ iff $\sigma(\tau) = \sigma(\tau')$
- $\mathcal{M} \models \varphi \wedge \psi$ iff $\mathcal{M} \models \varphi$ and $\mathcal{M} \models \psi$
- $\mathcal{M} \models \varphi \vee \psi$ iff $\mathcal{M} \models \varphi$ or $\mathcal{M} \models \psi$
- $\mathcal{M} \models \varphi \rightarrow \psi$ iff $\langle (D, \sigma), X, T \rangle \not\models \varphi$ or $\langle (D, \sigma), X, T \rangle \models \psi$, for $X \in \{H, T\}$
- $\mathcal{M} \models \forall x \varphi(x)$ iff $\mathcal{M} \models \varphi(d)$, for all $d \in D$
- $\mathcal{M} \models \exists x \varphi(x)$ iff $\mathcal{M} \models \varphi(d)$, for some $d \in D$ □

Equilibrium models for first-order theories are defined as follows.

Definition 4 (Quantified Equilibrium Model; Pearce and Valverde 2008)

Let φ be a first-order formula. A total QHT-interpretation $\mathcal{M} = \langle (D, \sigma), T, T \rangle$ is a first-order equilibrium model of φ if $\mathcal{M} \models \varphi$ and there is no model $\mathcal{M}' < \mathcal{M}$ of φ . □

We now focus on a particular fragment of QHT, called MHT(<), by imposing the following restrictions:

1. $C = [0.. \lambda)$ where $\lambda \in \mathbb{N}$ or $\lambda = \omega$ and $D = \{\mathbf{u}\} \cup C$ where \mathbf{u} stands for “undefined.”
2. We only allow for (unary) functions “+1” and “-1” with the expected meaning:

$$\sigma(\tau + 1) \stackrel{def}{=} \begin{cases} \sigma(\tau) + 1 & \text{if } \sigma(\tau) \neq \mathbf{u} \text{ and } \sigma(\tau) + 1 < \lambda \\ \mathbf{u} & \text{otherwise} \end{cases}$$

$$\sigma(\tau - 1) \stackrel{def}{=} \begin{cases} \sigma(\tau) - 1 & \text{if } \sigma(\tau) \neq \mathbf{u} \text{ and } \sigma(\tau) - 1 \geq 0 \\ \mathbf{u} & \text{otherwise} \end{cases}$$

3. All predicates are unary, except binary predicates = and <, interpreted as:

- (a) $\mathcal{M} \models \tau = \tau'$ if $\sigma(\tau) = \sigma(\tau') \neq \mathbf{u}$.
- (b) $\mathcal{M} \models \tau < \tau'$ if $\sigma(\tau) < \sigma(\tau')$ and both $\sigma(\tau) \neq \mathbf{u}$ and $\sigma(\tau') \neq \mathbf{u}$.

Note that the interpretation for equality requires now that both terms are different from \mathbf{u} . We define the abbreviation $x \leq y$ as $x < y \vee x = y$. Given these restrictions, we can simply represent an MHT($<$) interpretation as $\mathcal{M} = \langle \lambda, H, T \rangle$. Moreover, it is easy to see that we can establish a one-to-one mapping between the latter and an HT-trace $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$ with $\lambda = |\mathbf{M}|$ so that $H = \{a(i) \mid a \in H_i, i \in [0..\lambda]\}$ and $T = \{a(i) \mid a \in T_i, i \in [0..\lambda]\}$. When this happens, we say that \mathbf{M} and \mathcal{M} are *corresponding* interpretations.

Example 1

The HT-trace $\langle \mathbf{H}, \mathbf{T} \rangle$ with $\mathbf{H} = \{a\} \cdot \emptyset \cdot \{b\}$ and $\mathbf{T} = \{a, b\} \cdot \{a\} \cdot \{b\}$ corresponds to the MHT($<$) interpretation $\langle 3, H, T \rangle$ where $H = \{a(0), b(2)\}$ and $T = \{a(0), b(0), a(1), b(2)\}$.

We proceed to adapt now Kamp’s translation to our setting in the following way.

Definition 5 (Kamp’s translation)

Let φ be a temporal formula over \mathcal{A} . Kamp’s translation of φ for some time point $k \in \mathbb{N}$, denoted by $[\varphi]_k$, is defined as follows:

$$\begin{aligned}
 [\perp]_k &\stackrel{def}{=} \perp \\
 [a]_k &\stackrel{def}{=} a(k), \text{ with } a \in \mathcal{A} \\
 [\alpha \otimes \beta]_k &\stackrel{def}{=} [\alpha]_k \otimes [\beta]_k \text{ for any connective } \otimes \in \{\wedge, \vee, \rightarrow\} \\
 [\circ\alpha]_k &\stackrel{def}{=} \exists x (x = k + 1 \wedge [\alpha]_x) \\
 [\alpha \cup \beta]_k &\stackrel{def}{=} \exists x (k \leq x \wedge [\beta]_x \wedge \forall y (k \leq y \wedge y < x \rightarrow [\alpha]_y)) \\
 [\alpha \mathbb{R} \beta]_k &\stackrel{def}{=} \forall x (k \leq x \rightarrow [\beta]_x \vee \exists y (k \leq y \wedge y < x \wedge [\alpha]_y)) \\
 [\alpha \mathbb{W} \beta]_k &\stackrel{def}{=} \forall x (k \leq x \wedge \forall y (k \leq y \wedge y < x \rightarrow [\beta]_y) \rightarrow [\alpha]_x) \\
 [\bullet\alpha]_k &\stackrel{def}{=} \exists x (x = k - 1 \wedge [\alpha]_x) \\
 [\alpha \mathbf{S} \beta]_k &\stackrel{def}{=} \exists x (x \leq k \wedge [\beta]_x \wedge \forall y (x < y \wedge y \leq k \rightarrow [\alpha]_y)) \\
 [\alpha \mathbf{T} \beta]_k &\stackrel{def}{=} \forall x (x \leq k \rightarrow [\beta]_x \vee \exists y (x < y \wedge y \leq k \wedge [\alpha]_y))
 \end{aligned}$$

⊗

We now prove that, when considering the model correspondence between MHT($<$) and THT, Kamp’s translation is sound.

Theorem 1

Let φ be a temporal formula over \mathcal{A} , $\mathbf{M} = \langle \mathbf{H}, \mathbf{T} \rangle$ a THT-interpretation over \mathcal{A} and $\mathcal{M} = \langle (D, \sigma), H, T \rangle$ its corresponding MHT($<$)-interpretation.

Then, $\mathbf{M}, k \models \varphi$ in THT iff $\mathcal{M} \models [\varphi]_k$ in MHT($<$).

⊗

Corollary 1

A total THT-interpretation $\mathbf{M} = \langle \mathbf{T}, \mathbf{T} \rangle$ is a temporal equilibrium model of a temporal formula φ iff its corresponding MHT($<$)-interpretation $\mathcal{M} = \langle |\mathbf{T}|, T, T \rangle$ is an equilibrium model of $[\varphi]_0$.

⊗

The translation of derived operators can be simplified in MHT($<$) as follows:

$$\begin{aligned}
 [\mathbf{I}]_k &\equiv \neg \exists x (x = k - 1) \equiv k = 0 \\
 [\widehat{\bullet}\alpha]_k &\equiv \forall x (x = k - 1 \rightarrow [\alpha]_x) \equiv k = 0 \vee [\alpha]_{k-1} \\
 [\blacklozenge\alpha]_k &\equiv \exists x (x \leq k \wedge [\alpha]_x)
 \end{aligned}$$

$$\begin{aligned}
 [\blacksquare\alpha]_k &\equiv \forall x (x \leq k \rightarrow [\alpha]_x) \\
 [\mathbb{F}]_k &\equiv \neg\exists x (x = k + 1) \\
 [\widehat{\circ}\alpha]_k &\equiv \forall x (x = k + 1 \rightarrow [\alpha]_x) \\
 [\diamond\alpha]_k &\equiv \exists x (k \leq x \wedge [\alpha]_x) \\
 [\square\alpha]_k &\equiv \forall x (k \leq x \rightarrow [\alpha]_x)
 \end{aligned}$$

As an example, the translation of formula (2), viz. $\square(\neg a \rightarrow \circ a)$, for $k = 0$ amounts to

$$\begin{aligned}
 &\forall x (0 \leq x \rightarrow (\neg a(x) \rightarrow \exists y (y = x + 1 \wedge a(y)))) \\
 \equiv &\forall x (\neg a(x) \rightarrow \exists y (y = x + 1 \wedge a(y))),
 \end{aligned}$$

since in $\text{MHT}(<)$, $x \geq 0$ for any x . For infinite traces, the existence of some $y = x + 1$ is always guaranteed, and the formula above can be further simplified into

$$\forall x (\neg a(x) \rightarrow a(x + 1)),$$

which is just a first-order logic representation of the rule⁶ “ $\mathbf{a(X+1)} \text{ :- not } \mathbf{a(X)}$ ” in ASP. Similarly, it is not difficult to check that the translation of (1) amounts to:

$$\begin{aligned}
 &\forall x (0 \leq x \wedge \exists y (y = x - 1 \wedge \text{loaded}(y)) \wedge \neg\text{unloaded}(x) \rightarrow \text{loaded}(x)) \\
 \equiv &\forall x (0 < x \wedge \text{loaded}(x - 1) \wedge \neg\text{unloaded}(x) \rightarrow \text{loaded}(x)),
 \end{aligned}$$

that corresponds in ASP to the rule:

$$\text{loaded}(X) \text{ :- loaded}(X-1), \text{ not unloaded}(X), X > 0.$$

3 Foundations of Temporal Here-and-There

In this section, we explore some of the fundamental properties of THT, the monotonic basis of TEL. The importance of THT with respect to TEL is analogous to the relevance of HT for Equilibrium Logic and ASP. In particular, THT is a suitable framework to study the TEL-equivalence of two alternative representations. In what follows, we prove that THT-equivalence is a necessary and sufficient condition for *strong equivalence*, we provide several interesting equivalences in THT and we also present an alternative three-valued characterization of this logic. Besides, we explain how THT can be translated to LTL adding auxiliary atoms, something that allows reducing the strong equivalence problem to LTL-satisfiability checking. The section concludes with some results for THT for infinite traces, including properties about inter-definability of operators and an axiomatization.

3.1 Strong equivalence

As happens in ASP, given that TEL is a non-monotonic formalism, it may be the case that two different temporal formulas α and β share the same temporal equilibrium models but behave differently when a common context γ is added. For this reason, it is usual

⁶ Although variable X is unsafe, in a practical implementation, we would add a domain predicate $\text{time}(X)$ to specify that X is a time point.

to consider the notion of *strong equivalence* instead. Two temporal formulas α, β are *strongly equivalent* iff $\text{TEL}(\alpha \wedge \gamma) = \text{TEL}(\beta \wedge \gamma)$ for any arbitrary temporal formula γ . As expected, the THT-equivalence of $\alpha \equiv \beta$ is a sufficient condition for strong equivalence, since temporal equilibrium models are the result of a selection among THT-models.

Proposition 7 (Aguado et al. 2013; Cabalar et al. 2018)

If two temporal formulas α and β satisfy $\alpha \equiv \beta$ then they are strongly equivalent. ⊠

The interest of THT is that equivalence in that logic is also a *necessary* condition for strong equivalence, as we prove next.

Lemma 1

Let α and β be two LTL-equivalent formulas and let γ be the theory containing a formula $\beta \rightarrow \Box(a \vee \neg a)$ for each atom $a \in \mathcal{A}$.

Then, the following conditions are equivalent:

1. There exists some $\mathbf{H} < \mathbf{T}$ such that $\langle \mathbf{H}, \mathbf{T} \rangle \not\models \alpha \rightarrow \beta$;
2. \mathbf{T} is a TS-model of $\{\beta\} \cup \gamma$ but not a TS-model of $\{\alpha\} \cup \gamma$. ⊠

Proposition 8

If two temporal formulas α and β are strongly equivalent then $\alpha \equiv \beta$. ⊠

Proof

We prove that if $\alpha \not\equiv \beta$ then there is some context theory Γ for which $\{\alpha\} \cup \Gamma$ and $\{\beta\} \cup \Gamma$ have different TS-models. Assume first that α and β have different total models, that is, different LTL-models. Then, take the set $\Gamma = (\text{EM})$ of excluded middle axioms for every $a \in \mathcal{A}$. The LTL-models of $\{\alpha\} \cup (\text{EM})$ and $\{\beta\} \cup (\text{EM})$ also differ (since (EM) is a set of LTL tautologies). But by Proposition 1, LTL-models of these theories are exactly their TS-models, and so, they also differ.

Suppose now that α and β are LTL-equivalent but still, $\alpha \not\equiv \beta$. Then, there is some THT-countermodel $\langle \mathbf{H}, \mathbf{T} \rangle$ of either $(\alpha \rightarrow \beta)$ or $(\beta \rightarrow \alpha)$, and given LTL-equivalence of α and β , the countermodel is non-total, $\mathbf{H} < \mathbf{T}$. Without loss of generality, assume $\langle \mathbf{H}, \mathbf{T} \rangle \not\models \alpha \rightarrow \beta$. By Lemma 1, taking the theory Γ consisting of an implication $\beta \rightarrow \Box(a \vee \neg a)$ for each atom $a \in \mathcal{A}$, we get that \mathbf{T} is TS-model of $\{\beta\} \cup \Gamma$ but not TS-model of $\{\alpha\} \cup \Gamma$. □

As a consequence of Propositions 7 and 8, we obtain the following characterization:

Theorem 2 (Strong Equivalence characterization)

Two temporal formulas α and β are strongly equivalent iff $\alpha \equiv \beta$. ⊠

3.2 Some interesting properties of THT

THT-equivalences have a crucial role for deciding how to rewrite a formula without caring about the possible context in which it is included. We analyze next several useful THT-equivalences. The following are some De Morgan laws satisfied by negation and other operators:

$$\neg(\varphi \wedge \psi) \equiv \neg\varphi \vee \neg\psi, \tag{3}$$

$$\neg(\varphi \vee \psi) \equiv \neg\varphi \wedge \neg\psi, \tag{4}$$

$$\neg(\varphi \cup \psi) \equiv \neg\varphi \mathbb{R} \neg\psi, \tag{5}$$

$$\neg(\varphi \mathbb{R} \psi) \equiv \neg\varphi \cup \neg\psi, \tag{6}$$

$$\neg(\varphi \mathbb{W} \psi) \equiv \neg\neg\psi \cup \neg\varphi, \tag{7}$$

$$\neg(\circ\varphi) \equiv \widehat{\circ}\neg\varphi, \tag{8}$$

$$\neg(\widehat{\circ}\varphi) \equiv \circ\neg\varphi. \tag{9}$$

The next operators distribute over disjunction and conjunction so that:

$$\circ(\varphi \oplus \psi) \equiv \circ\varphi \oplus \circ\psi, \tag{10}$$

$$\widehat{\circ}(\varphi \oplus \psi) \equiv \widehat{\circ}\varphi \oplus \widehat{\circ}\psi, \tag{11}$$

for $\oplus \in \{\vee, \wedge\}$ but for implication and temporal operators, only these distributive equivalences are valid:

$$\widehat{\circ}(\varphi \rightarrow \psi) \equiv \widehat{\circ}\varphi \rightarrow \widehat{\circ}\psi, \tag{12}$$

$$\circ(\varphi \cup \psi) \equiv \circ\varphi \cup \circ\psi, \tag{13}$$

$$\circ\Diamond\varphi \equiv \Diamond\circ\varphi, \tag{14}$$

$$\widehat{\circ}(\varphi \mathbb{R} \psi) \equiv \widehat{\circ}\varphi \mathbb{R} \widehat{\circ}\psi, \tag{15}$$

$$\widehat{\circ}(\varphi \mathbb{W} \psi) \equiv \widehat{\circ}\varphi \mathbb{W} \widehat{\circ}\psi, \tag{16}$$

$$\widehat{\circ}\Box\varphi \equiv \Box\widehat{\circ}\varphi. \tag{17}$$

All these properties are symmetrically satisfied by the past-oriented versions of the operators above. For infinite traces, however, we obtain the additional relation:

$$\widehat{\circ}\varphi \equiv \circ\varphi, \tag{18}$$

so both next operators coincide. This equivalence never holds for previous operators, \bullet and $\widehat{\bullet}$, since there always exists an initial state.

Also, THT satisfies *persistence*, a characteristic property of HT-based logics.

Proposition 9 (Persistence; Aguado et al. 2013; Cabalar et al. 2018)

Let $\langle \mathbf{H}, \mathbf{T} \rangle$ be an HT-trace of length λ and φ be a temporal formula.

Then, for any $k \in [0..\lambda)$, if $\langle \mathbf{H}, \mathbf{T} \rangle, k \models \varphi$ then $\langle \mathbf{T}, \mathbf{T} \rangle, k \models \varphi$ (or, if preferred, $\mathbf{T}, k \models \varphi$). ⊠

As a corollary, we have that $\langle \mathbf{H}, \mathbf{T} \rangle \models \neg\varphi$ iff $\mathbf{T} \not\models \varphi$ in LTL. All THT tautologies are LTL tautologies but not vice versa. However, they coincide for some types of equivalences, as we show next.

Proposition 10 (Aguado et al. 2013; Cabalar et al. 2018)

Let φ and ψ be temporal formulas without implications (and so, without negations either).

Then, $\varphi \equiv \psi$ in LTL iff $\varphi \equiv \psi$ in THT. ⊠

As an example, the usual inductive definition of the until operator from LTL

$$\varphi \cup \psi \equiv \psi \vee (\varphi \wedge \circ(\varphi \cup \psi)), \tag{19}$$

is also valid in THT due to Proposition 10. In fact, by De Morgan laws, LTL satisfies a kind of duality guaranteeing, for instance, that (19) iff

$$\varphi \mathbb{R} \psi \equiv \psi \wedge (\varphi \vee \widehat{\circ}(\varphi \mathbb{R} \psi)), \tag{20}$$

and, by Proposition 10 again, this is also a valid equivalence in THT.

If we define all the pairs of dual connectives as follows: \wedge/\vee , \top/\perp , \mathbb{U}/\mathbb{R} , $\circ/\widehat{\circ}$, \square/\diamond , \mathbf{S}/\mathbf{T} , $\bullet/\widehat{\bullet}$, $\blacksquare/\blacklozenge$, we can extend this to any formula φ without implications and define $\delta(\varphi)$ as the result of replacing each connective by its dual operator. Then, we get the following corollary of Proposition 10.

Corollary 2 (Boolean Duality; Cabalar et al. 2018)

Let φ and ψ be formulas without implication.

Then, THT satisfies: $\varphi \equiv \psi$ iff $\delta(\varphi) \equiv \delta(\psi)$. ⊠

In a similar manner, the temporal symmetry in the system can be exploited in order to switch the temporal direction of operators to conclude, for instance, that (19) iff $\varphi \mathbf{S} \psi \equiv \psi \vee (\varphi \wedge \bullet(\varphi \mathbf{S} \psi))$. However, this duality has some obvious limitations when we allow for infinite traces. For instance, the past has a beginning $\blacklozenge \mathbf{I} \equiv \top$ but the future may have no end $\diamond \mathbb{F} \not\equiv \top$. If we restrict ourselves to finite traces, we get the following result.

To this end, let \mathbb{U}/\mathbf{S} , \mathbb{R}/\mathbf{T} , \circ/\bullet , $\widehat{\circ}/\widehat{\bullet}$, \square/\blacksquare , and \diamond/\blacklozenge denote all pairs of swapped-time connectives and let $\sigma(\varphi)$ denote the replacement in φ of each connective by its swapped-time version.

Lemma 2 (Cabalar et al. 2018)

There exists a mapping ϱ on finite HT-traces of the same length λ such that for any $k \in [0..\lambda)$, $\mathbf{M}, k \models \varphi$ iff $\varrho(\mathbf{M}), \lambda - 1 - k \models \sigma(\varphi)$. ⊠

Theorem 3 (Temporal Duality Theorem; Cabalar et al. 2018)

A temporal formula φ is a THT_f-tautology iff $\sigma(\varphi)$ is a THT_f-tautology. ⊠

3.3 Three-valued characterization of THT

As mentioned in Section 2, HT-traces can be seen as three-valued models where each atom can be false, assumed or proven (in a given state). This stems from the fact that the intermediate logic of HT actually corresponds to Gödel’s three-valued logic G_3 (Gödel 1932). In particular, the HT-satisfaction of formulas can be replaced by a three-valued function that assigns, to each formula, a value 0, 1, or 2, standing for “false” (it does not hold in T), “assumed” (it holds in T but not in H) and “proven” (it holds in H , and so, in T , too), respectively. The interest of a multi-valued truth assignment $\mathbf{m}(k, \alpha)$ of a formula α at time point k is that THT-equivalence $\alpha \equiv \beta$ can be reduced to a comparison of truth values such as $\mathbf{m}(k, \alpha) = \mathbf{m}(k, \beta)$. This has an important application when introducing an auxiliary atom a to represent α so that we can safely replace α by a provided that we add formulas to guarantee $\mathbf{m}(k, a) = \mathbf{m}(k, \alpha)$. Following these ideas, we proceed with the following formal definitions.

For convenience, the evaluation of implications is captured by the following conditional operation on truth values: $\text{imp}(x, y) = 2$ if $x \leq y$; otherwise $\text{imp}(x, y) = y$. Given an HT-trace $\langle \mathbf{H}, \mathbf{T} \rangle$ of length λ , we define its associated truth valuation as a function $\mathbf{m}(k, \varphi)$ that assigns a truth value in $\{0, 1, 2\}$ to formula φ at time point $k \in [0..\lambda)$ according to the following rules:

$$\begin{aligned}
 \mathbf{m}(k, \perp) &\stackrel{def}{=} 0 \\
 \mathbf{m}(k, a) &\stackrel{def}{=} \begin{cases} 0 & \text{if } a \notin T_k \\ 1 & \text{if } a \in T_k \setminus H_k \\ 2 & \text{if } a \in H_k \end{cases} \quad \text{for any atom } a \\
 \mathbf{m}(k, \varphi \wedge \psi) &\stackrel{def}{=} \min(\mathbf{m}(k, \varphi), \mathbf{m}(k, \psi)) \\
 \mathbf{m}(k, \varphi \vee \psi) &\stackrel{def}{=} \max(\mathbf{m}(k, \varphi), \mathbf{m}(k, \psi)) \\
 \mathbf{m}(k, \varphi \rightarrow \psi) &\stackrel{def}{=} \text{imp}(\mathbf{m}(k, \varphi), \mathbf{m}(k, \psi)) \\
 \mathbf{m}(k, \bullet\varphi) &\stackrel{def}{=} \begin{cases} 0 & \text{if } k = 0 \\ \mathbf{m}(k - 1, \varphi) & \text{if } k > 0 \end{cases} \\
 \mathbf{m}(k, \varphi \mathbf{S} \psi) &\stackrel{def}{=} \max\{\min(\mathbf{m}(j, \psi), \min\{\mathbf{m}(i, \varphi) \mid j < i \leq k\}) \mid 0 \leq j \leq k\} \\
 \mathbf{m}(k, \varphi \mathbf{T} \psi) &\stackrel{def}{=} \min\{\max(\mathbf{m}(j, \psi), \max\{\mathbf{m}(i, \varphi) \mid j < i \leq k\}) \mid 0 \leq j \leq k\} \\
 \mathbf{m}(k, \circ\varphi) &\stackrel{def}{=} \begin{cases} 0 & \text{if } k + 1 = \lambda (\neq \omega) \\ \mathbf{m}(k + 1, \varphi) & \text{if } k + 1 < \lambda \end{cases} \\
 \mathbf{m}(k, \varphi \mathbb{U} \psi) &\stackrel{def}{=} \max\{\min(\mathbf{m}(j, \psi), \min\{\mathbf{m}(i, \varphi) \mid k \leq i < j\}) \mid k \leq j < \lambda\} \\
 \mathbf{m}(k, \varphi \mathbb{R} \psi) &\stackrel{def}{=} \min\{\max(\mathbf{m}(j, \psi), \max\{\mathbf{m}(i, \varphi) \mid k \leq i < j\}) \mid k \leq j < \lambda\} \\
 \mathbf{m}(k, \varphi \mathbb{W} \psi) &\stackrel{def}{=} \min\{\text{imp}(\min\{\mathbf{m}(i, \psi) \mid k \leq i < j\}, \mathbf{m}(j, \varphi)) \mid k \leq j < \lambda\}
 \end{aligned}$$

The valuation of derived operators can be easily concluded from their definitions:

$$\begin{aligned}
 \mathbf{m}(k, \top) &= 2 \\
 \mathbf{m}(k, \neg\varphi) &\stackrel{def}{=} \begin{cases} 2 & \text{if } \mathbf{m}(k, \varphi) = 0 \\ 0 & \text{otherwise} \end{cases} \\
 \mathbf{m}(k, \mathbf{1}) &\stackrel{def}{=} \begin{cases} 2 & \text{if } k = 0 \\ 0 & \text{if } k > 0 \end{cases} \\
 \mathbf{m}(k, \widehat{\bullet}\varphi) &\stackrel{def}{=} \begin{cases} 2 & \text{if } k = 0 \\ \mathbf{m}(k - 1, \varphi) & \text{if } k > 0 \end{cases} \\
 \mathbf{m}(k, \blacksquare\varphi) &= \min\{\mathbf{m}(i, \varphi) \mid 0 \leq i \leq k\} \\
 \mathbf{m}(k, \blacklozenge\varphi) &= \max\{\mathbf{m}(i, \varphi) \mid 0 \leq i \leq k\} \\
 \mathbf{m}(k, \mathbb{F}) &\stackrel{def}{=} \begin{cases} 2 & \text{if } k = \lambda \neq \omega \\ 0 & \text{if } k < \lambda \end{cases} \\
 \mathbf{m}(k, \widehat{\circ}\varphi) &\stackrel{def}{=} \begin{cases} 2 & \text{if } k + 1 = \lambda (\neq \omega) \\ \mathbf{m}(k + 1, \varphi) & \text{if } k + 1 < \lambda \end{cases} \\
 \mathbf{m}(k, \square\varphi) &= \min\{\mathbf{m}(i, \varphi) \mid k \leq i < \lambda\} \\
 \mathbf{m}(k, \diamond\varphi) &= \max\{\mathbf{m}(i, \varphi) \mid k \leq i < \lambda\}
 \end{aligned}$$

For the restriction to THT_f , it suffices to impose the condition $\lambda \neq \omega$.

Proposition 11

Let $\langle \mathbf{H}, \mathbf{T} \rangle$ be a HT-trace of length λ , \mathbf{m} its associated valuation and $k \in [0.. \lambda)$.

Then, we have that

- $\langle \mathbf{H}, \mathbf{T} \rangle, k \models \varphi$ iff $\mathbf{m}(k, \varphi) = 2$
- $\langle \mathbf{T}, \mathbf{T} \rangle, k \models \varphi$ iff $\mathbf{m}(k, \varphi) \neq 0$ ⊠

As an illustration, consider the HT-trace from Example 1. This trace corresponds to the three-valued assignment

$$\mathbf{m}(0, a) = 2, \mathbf{m}(0, b) = 1, \mathbf{m}(1, a) = 1, \mathbf{m}(1, b) = 0, \mathbf{m}(2, a) = 0 \text{ and } \mathbf{m}(2, b) = 2.$$

The formula $\diamond b$ is proven at $k = 0$, that is, $\mathbf{m}(0, \diamond b) = 2$, because $\mathbf{m}(0, b) = 2$ and 2 is the maximum possible value. The formula $\square b$ is false at $k = 0$, $\mathbf{m}(0, \square b) = 0$ because $\mathbf{m}(1, b) = 0$ and 0 is the minimum value. On the other hand, the formula $\square(a \vee b)$ is just assumed but not proven at 0. To see why, note that the value $\mathbf{m}(k, a \vee b)$ is the maximum of both disjuncts and that, in our trace, this is equal to 2 for time points $k = 0$ and $k = 2$ but is equal to 1 for $k = 1$. Since $\square(a \vee b)$ takes the minimum of all these values for $k \geq 0$, we get $\mathbf{m}(0, \square(a \vee b)) = 1$.

3.4 From THT to LTL

Propositions 1 and 2 tell us that the addition of excluded middle axioms (EM) allows for an easy encoding of LTL into THT. In fact, Proposition 6 even guarantees that the addition of these axioms makes also the non-monotonic formalism of TEL collapse into LTL. An interesting question is whether the other direction is possible, namely, whether THT can be encoded into LTL. In fact, this can be done by adapting the encoding of (non-temporal) HT into classical propositional logic, as presented, for instance, by Pearce et al. (2001).

Given a propositional signature \mathcal{A} , let us define a new propositional signature in LTL by $\mathcal{A}^* = \mathcal{A} \cup \{a' \mid a \in \mathcal{A}\}$. For any temporal formula φ , we define its translation φ^* as follows:

1. $\perp^* \stackrel{def}{=} \perp$
2. $a^* \stackrel{def}{=} a'$ for any $a \in \mathcal{A}$
3. $(\odot \varphi)^* \stackrel{def}{=} \odot \varphi^*$, if $\odot \in \{\circ, \bullet\}$
4. $(\varphi \odot \psi)^* \stackrel{def}{=} \varphi^* \odot \psi^*$, when $\odot \in \{\wedge, \vee, \sqcup, \mathbb{R}, \mathbf{S}, \mathbf{T}\}$
5. $(\varphi \rightarrow \psi)^* \stackrel{def}{=} (\varphi \rightarrow \psi) \wedge (\varphi^* \rightarrow \psi^*)$

We associate with any HT-trace $\langle \mathbf{H}, \mathbf{T} \rangle$ of length λ the trace \mathbf{T}^* in LTL defined as the sequence of sets of atoms $T_i^* = T_i \cup \{a' \mid a \in H_i\}$ for any $i \in [0..\lambda)$. Informally speaking, \mathbf{T}^* considers a new primed atom a' per each $a \in \mathcal{A}$ in the original signature. In the trace, the primed atom a' represents the fact that a occurs at some point in the \mathbf{H} components, whereas the original symbol a is used to represent an atom in \mathbf{T} . As an HT-trace satisfies $H_i \subseteq T_i$ by construction, we may have traces that do not correspond to any HT-trace, since the set of primed atoms must be a subset of the non-primed ones. To force this structural condition on traces, we can just add the axiom:

$$\square(a' \rightarrow a) \text{ for any atom } a \in \mathcal{A}. \tag{21}$$

By including this axiom, we obtain a one-to-one correspondence between HT-traces and traces for the extended signature. In particular, if we take any arbitrary trace \mathbf{T}^* for

signature \mathcal{A}^* satisfying (21), we can now define its corresponding HT-trace $\langle \mathbf{H}, \mathbf{T} \rangle$ as $T_i \stackrel{def}{=} T_i^* \cap \mathcal{A}$ and $H_i \stackrel{def}{=} \{a \mid a' \in \mathbf{T}_i^*\}$.

Theorem 4

Let $\langle \mathbf{H}, \mathbf{T} \rangle$ be an HT-trace of length λ for alphabet \mathcal{A} , $k \in [0..\lambda)$ and φ a temporal formula over \mathcal{A} . Then, $\langle \mathbf{H}, \mathbf{T} \rangle, k \models \varphi$ iff $\mathbf{T}^*, k \models \varphi^* \wedge$ (21) in LTL. \square

Example 2

Take the HT-trace $\langle \mathbf{H}, \mathbf{T} \rangle$ from Example 1. Its corresponding trace is $\mathbf{T}^* = \{a, b, a'\} \cdot \{a\} \cdot \{b, b'\}$. Axiom (21) in this case is the formula:

$$\square(a' \rightarrow a) \wedge \square(b' \rightarrow b). \tag{22}$$

Now, for instance, given the formula:

$$\square(\neg b \rightarrow a), \tag{23}$$

its translation (23)* corresponds to:

$$\begin{aligned} & \square((\neg b \rightarrow a) \wedge ((\neg b)^* \rightarrow a^*)) \\ = & \square((\neg b \rightarrow a) \wedge (\neg b \wedge \neg b' \rightarrow a')) \\ \equiv & \square(\neg b \rightarrow a) \wedge \square(\neg b \rightarrow a'). \end{aligned} \tag{24}$$

It is easy to see that \mathbf{T}^* satisfies (23)* \wedge (22) in LTL.

To conclude this comparison to LTL, we include a brief comment on the complexity of THT. As shown above, LTL-satisfaction of a temporal formula φ can be reduced to THT-satisfaction with the simple addition of (EM) axioms. On the other hand, reducing THT-satisfaction of φ to LTL-satisfaction consists of adding axiom (21) and applying translation φ^* , which can be easily proved to be quadratic in the worse case. As a result, LTL- and THT-satisfaction share the same complexity, regardless of the trace length under consideration. In particular, it is well-known that LTL_ω -satisfiability is PSPACE-complete (Sistla and Clarke 1985). Also, LTL_f -satisfiability was also proved to be PSPACE-complete by De Giacomo and Vardi (2013). As a result, both THT_ω - and THT_f -satisfiability have the same complexity, too.

3.5 Automata-based checking of strong equivalence

The translation from THT to LTL can be used to reduce strong equivalence of temporal theories to LTL-satisfiability checking. This feature was exploited by the tool **abstem** (Cabalar and Diéguez 2014) that provides us with several functionalities for temporal theories under TEL_ω semantics.

For checking whether two temporal formulas φ and ψ are LTL-equivalent, **abstem** checks the validity of $\varphi \leftrightarrow \psi$, which is reduced to the unsatisfiability of $\neg(\varphi \leftrightarrow \psi)$. This is done by translating $\neg(\varphi \leftrightarrow \psi)$ into a Büchi automaton \mathfrak{A} by means of the LTL-model checker **SPoT** (Duret-Lutz et al. 2016). As explained in Section 4, the language accepted by \mathfrak{A} corresponds to the LTL-models of $\neg(\varphi \leftrightarrow \psi)$. Therefore, if the accepting language of \mathfrak{A} is empty, it means that $\neg(\varphi \leftrightarrow \psi)$ is not satisfiable and φ and ψ are LTL-equivalent. Otherwise, any word accepted by \mathfrak{A} can be seen as a counterexample of $\varphi \leftrightarrow \psi$.

The method used by Cabalar and Diéguez (2014) to determine whether φ and ψ are strongly equivalent consists in checking the validity of both $\varphi \rightarrow \psi$ and $\psi \rightarrow \varphi$ in THT. Let us consider, without loss of generality, that we want to check the validity of $\varphi \rightarrow \psi$. This is equivalent to checking the satisfiability of

$$\neg \left(\left(\bigwedge_{a \in \mathcal{A}} \Box (a' \rightarrow a) \right) \rightarrow (\varphi \rightarrow \psi)^* \right), \tag{25}$$

in LTL and it can be done by using the procedure explained above. `abstem` obtains a Büchi automaton $\mathfrak{A}_{(25)}$ that accepts the LTL-models of (25). If (25) is LTL-satisfiable then `abstem` filters all the variables of the type a' from the accepting language of $\mathfrak{A}_{(25)}$ (as explained in Section 4). The resulting automaton, denoted by $h(\mathfrak{A}_{(25)})$, captures the TEL-models of $\varphi \wedge \gamma$ which are not models of $\psi \wedge \gamma$, where

$$\gamma \stackrel{def}{=} \bigwedge_{a \in \mathcal{A}} \psi \rightarrow \Box(a \vee \neg a),$$

for every atom a in the signature (assuming a finite alphabet).

With minor modifications, this technique could be adapted to the TEL_f case. The only required change is the target automata: for TEL_ω we need to use Büchi automata since we are dealing with infinite computations while, in the finite case, we need to use automata on finite words. The rest of the algorithm would remain the same.

3.6 Definability of temporal operators in THT

Propositional connectives are not inter-definable in intuitionistic propositional logic (van Dalen 2001). However, when it comes to propositional HT, some inter-definability results can be obtained. For instance, the HT-valid formula (Lukasiewicz 1941)

$$p \vee q \leftrightarrow ((p \rightarrow q) \rightarrow q) \wedge ((q \rightarrow p) \rightarrow p).$$

allows us to determine that disjunction in HT can be defined in terms of the remaining propositional connectives. Unfortunately, this is the only definable operator (Balbiani and Diéguez 2016; Aguado et al. 2015). When considering QHT, the equivalence

$$\exists x p(x) \leftrightarrow (\forall y \forall x (p(x) \rightarrow p(y)) \rightarrow p(y)),$$

is valid, so existential quantifiers can be reformulated in terms of the remaining connectives as well (Mints 2010). Since THT can be seen as a syntactic subclass of QHT, a similar result could be expected. In fact, Balbiani et al. (2020) have proved that the connectives $\diamond p$ can be defined in terms of a \cup -free formula, as stated in the following lemma.

Lemma 3 (Balbiani et al. 2020)

The formulas $\diamond p$ and

$$(\Box(p \rightarrow \Box(p \vee \neg p)) \wedge \Box(\Box(p \vee \neg p) \rightarrow p \vee \neg p \vee \Box \neg p)) \rightarrow (\Box(p \vee \neg p) \wedge \neg \Box \neg p), \tag{26}$$

are equivalent in THT.

Broadly speaking the equivalence in Lemma 3 captures the three different ways of satisfying $\diamond p$ on any THT-model $\langle \mathbf{H}, \mathbf{T} \rangle$ at $i \geq 0$:

1. $\langle \mathbf{H}, \mathbf{T} \rangle, i \models \Box(p \vee \neg p)$ holds, and in this case $\langle \mathbf{H}, \mathbf{T} \rangle$ behaves classically after i ; at least for formulas whose only variable is p . In this case $\diamond p$ would behave as $\neg \Box \neg p$.

2. If $\langle \mathbf{H}, \mathbf{T} \rangle, i \not\models \Box(p \vee \neg p)$, then $\langle \mathbf{H}, \mathbf{T} \rangle$ does not behave classically after i . Therefore, for some $j \geq i$, $\langle \mathbf{H}, \mathbf{T} \rangle, i \not\models p \vee \neg p$. In order to make the right part of the equivalence true, either $\langle \mathbf{H}, \mathbf{T} \rangle, i \not\models \Box(p \rightarrow \Box(p \vee \neg p))$ or $\langle \mathbf{H}, \mathbf{T} \rangle, i \not\models \Box(\Box(p \vee \neg p) \rightarrow p \vee \neg p \vee \Box \Box \neg p)$ holds. The former formula fails when there is $k \geq i$ such that $\langle \mathbf{H}, \mathbf{T} \rangle, j \models p$ (so $\langle \mathbf{H}, \mathbf{T} \rangle, i \models \Diamond p$) and $\langle \mathbf{H}, \mathbf{T} \rangle$ does not behave classically after k .
3. Similarly, $\Box(\Box(p \vee \neg p) \rightarrow p \vee \neg p \vee \Box \Box \neg p)$ fails exactly where there is $k \geq i$ satisfying p but $\langle \mathbf{H}, \mathbf{T} \rangle$ behaves classically after k . In other words $\langle \mathbf{H}, \mathbf{T} \rangle$ falsifies $p \vee \neg p$ only for $i < k$. In this case, $\Box(p \vee \neg p) \rightarrow p \vee \neg p \vee \Box \Box \neg p$ is falsified exactly at the greatest such k .

From this equivalence the following corollary follows.

Corollary 3 (Balbiani et al. 2020)

$p \cup q$ is \cup -free definable by using the equivalence $q \cup p \leftrightarrow (p \mathbb{R} (q \vee p)) \wedge (26)$.

Lemma 3 and Corollary 3 were proved for infinite traces, but in fact, these equivalences also hold for finite traces THT_f , as stated below.

Corollary 4

The equivalences (26) and $q \cup p \leftrightarrow (p \mathbb{R} (q \vee p)) \wedge (26)$ are valid in THT_f .

Balbiani et al. (2020) also give a negative answer for the definability of $\Box p$ or $p \mathbb{R} q$ as a \mathbb{R} -free formula. To do so, they consider the following THT model: let $n \in \mathbb{N}$, we define the THT models $\langle \mathbf{H}, \mathbf{T} \rangle_n$ as follows: (1) $H_i = \{p\}$ for all $0 \leq i < n + 1$; (2) $H_{n+1} = \emptyset$; (3) $T_i = \{p\}$, for all $0 \leq i \leq n + 1$; (4) $T_i = T_{i \bmod n+2}$ and $H_i = H_{i \bmod n+2}$, for all $i > n + 1$; It is evident that $\langle \mathbf{H}, \mathbf{T} \rangle, 0 \not\models \Box p$ and $\langle \mathbf{T}, \mathbf{T} \rangle, 0 \models \Box p$. By using a *bisimulation* argument, Balbiani et al. (2020) proved the following lemma.

Lemma 4 (Balbiani et al. 2020)

Let us consider the THT model defined above. The models $\langle \mathbf{H}, \mathbf{T} \rangle_n$ and $\langle \mathbf{T}, \mathbf{T} \rangle_n$ satisfy the same \mathbb{R} -free formulas with at most n connectives at time 0.

This lemma is proved by showing that $\langle \mathbf{H}, \mathbf{T} \rangle_n$ and $\langle \mathbf{T}, \mathbf{T} \rangle_n$ are *bisimilar* at time 0 with respect to the temporal language excluding the \mathbb{R} connective. This means both models cannot be distinguish with respect to the set of \mathbb{R} -free formulas.

As a consequence, let us assume by contradiction that $\Box p$ is definable in terms of a \mathbb{R} -free formula ψ and let us fix n as the number of connectives of ψ . Since $\langle \mathbf{H}, \mathbf{T} \rangle_n, 0 \not\models \Box p$ then $\langle \mathbf{H}, \mathbf{T} \rangle_n, 0 \not\models \psi$. Since $\langle \mathbf{T}, \mathbf{T} \rangle_n, 0 \models \Box p$ then $\langle \mathbf{T}, \mathbf{T} \rangle_n, 0 \models \psi$, which contradicts Lemma 4. As a corollary, if we could define the formula $p \mathbb{R} q$ in terms of a \mathbb{R} -free formula, we could also define $\Box p$, since the latter is equivalent to $p \mathbb{R} \perp$.

If we remove condition (4) from the model $\langle \mathbf{H}, \mathbf{T} \rangle_n$ we obtain a finite model with $\lambda = n + 2$ satisfying Lemma 4. Therefore, we can reformulate the previous discussion in terms of THT_f .

Another interesting observation regarding expressiveness of operators is the presence of the newly introduced temporal operator $\varphi \mathbb{W} \psi$ read as “repeat φ while ψ .” Another interesting observation regarding expressiveness of operators is the presence of the newly introduced temporal operator $\varphi \mathbb{W} \psi$, read as “repeat φ while ψ .” This operator is not usual in LTL because, in that logic, it can be easily defined in terms of \cup or \mathbb{R} :

Proposition 12

In LTL, we have the following equivalences:

$$\begin{aligned} \alpha \mathbb{W} \beta &\equiv \neg(\beta \cup \neg\alpha) \equiv \neg\beta \mathbb{R} \alpha \\ \alpha \cup \beta &\equiv \neg\neg(\alpha \cup \neg\neg\beta) \\ &\equiv \neg(\neg\beta \mathbb{W} \alpha) \end{aligned}$$

⊠

In THT, however, these equivalences do not hold. Unlike \cup and \mathbb{R} , operator \mathbb{W} has an implicational nature. As a consequence, there is no clear way to represent \mathbb{W} in terms of the other operators: we conjecture that this may be, in fact, a fundamental operator in THT. The following equivalences describe these three operators, \cup , \mathbb{R} and \mathbb{W} , by an inductive unfolding.

Proposition 13

In THT, we have the following equivalences:

$$\begin{aligned} \alpha \cup \beta &\equiv \beta \vee (\alpha \wedge \circ(\alpha \cup \beta)) \\ \alpha \mathbb{R} \beta &\equiv \beta \wedge (\alpha \vee \circ(\alpha \mathbb{R} \beta)) \\ \alpha \mathbb{W} \beta &\equiv \alpha \wedge (\beta \rightarrow \circ(\alpha \mathbb{W} \beta)) \end{aligned}$$

⊠

The first two equivalences are well-known from LTL and also preserved in THT. The third equivalence shows that the \mathbb{W} operator is formed by a repeated application of implications. This is even clearer in the following expansion.

Proposition 14

The formula $(\alpha \mathbb{W} \beta)$ is THT-equivalent to the (possibly infinite) conjunction of rules

$$(\bigwedge_{j=0}^{i-1} \circ^j \beta) \rightarrow \circ^i \alpha \quad \text{for all } i \geq 0.$$

⊠

To see how this expression works, consider the formula $(w \mathbb{W} f)$, where w means “pouring water” and f means “fire.” Notice that, for $i = 0$, the rule body becomes an empty conjunction (\top) and so this expands to $\top \rightarrow \circ^0 w$ which is just equivalent to fact w . For instance, for $i \in [0, 3]$ we get the rules:

$$\begin{aligned} &w \\ f &\rightarrow \circ w \\ f \wedge \circ f &\rightarrow \circ^2 w \\ f \wedge \circ f \wedge \circ^2 f &\rightarrow \circ^3 w \\ &\vdots \end{aligned}$$

That is, we start pouring water at situation 0 and, if we have fire, we keep pouring water at 1 and check fire again, and so on. As we can see, the effect of each f test is placed at the next situation. Thus, the reading of $(w \mathbb{W} f)$ is like a procedural program “**do** pour-water **while** fire”. If we want to move the test to the same situation of its effect, we

would write instead $((f \rightarrow w) \mathbb{W} f)$ whose reading would be “**while** fire **do** pour-water” and whose expansion becomes:

$$\begin{aligned} f &\rightarrow w \\ f \wedge \circ f &\rightarrow \circ w \\ f \wedge \circ f \wedge \circ^2 f &\rightarrow \circ^2 w \\ &\vdots \end{aligned}$$

The expansion of $(w \mathbb{W} f)$ as a set of rules reveals its behavior from a logic programming point of view. For instance, a theory only containing the formula $(w \mathbb{W} f)$ has a unique temporal stable model (per each length $\lambda > 0$) in which w is only true at the initial state, whereas f is always false, since there is no additional evidence about fire. In LTL, the formula $(w \mathbb{W} f)$ is equivalent to $(\neg f \mathbb{R} w)$, that is, $\Box w \vee (w \mathbb{U} (w \wedge \neg f))$. In our formalism, however, this last formula produces temporal stable models with arbitrary prefix sequences of w , and even the case in which w holds in all the states of the trace. In other words, water may be poured arbitrarily many times, even though fire is false all over the trace.

3.7 Axiomatization of THT_ω

The axiomatic system of the future fragment of THT_ω (Balbiani and Diéguez 2016) is obtained by combining the axioms of intuitionistic modal logic K (Simpson 1994) (IK), the Hosiøi Axiom (Hosiøi 1966) and the future LTL $_\omega$ axiomatization of Goldblatt (1992). Such an axiomatic system is described below.

1. Axioms of Intuitionistic Propositional Calculus

2. **Hosiøi axiom:** $p \vee (p \rightarrow q) \vee \neg q$

3. **IK axioms for \circ and $\widehat{\circ}$** (Simpson 1994):

(a) $\widehat{\circ}p \leftrightarrow \circ p$	(d) $\circ(p \vee q) \leftrightarrow \circ p \vee \circ q$
(b) $\widehat{\circ}(p \rightarrow q) \rightarrow (\widehat{\circ}p \rightarrow \widehat{\circ}q)$	(e) $(\circ p \rightarrow \widehat{\circ}q) \rightarrow \widehat{\circ}(p \rightarrow q)$
(c) $\widehat{\circ}(p \rightarrow q) \rightarrow (\circ p \rightarrow \circ q)$	(f) $\neg \circ \perp$

4. IK axioms for \Box and \Diamond :

(g) $\Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)$	(j) $(\Diamond p \rightarrow \Box q) \rightarrow \Box(p \rightarrow q)$
(h) $\Box(p \rightarrow q) \rightarrow (\Diamond p \rightarrow \Diamond q)$	(k) $\neg \Diamond \perp$
(i) $\Diamond(p \vee q) \rightarrow \Diamond p \vee \Diamond q$	

5. Axioms combining \circ , $\widehat{\circ}$, \Box and \Diamond :

(l) $\Box p \rightarrow p \wedge \widehat{\circ}\Box p$	(m) $p \vee \circ \Diamond p \rightarrow \Diamond p$
---	--

6. Induction:

(n) $\frac{p \rightarrow \widehat{\circ}p}{p \rightarrow \Box p}$	(o) $\frac{\circ p \rightarrow p}{\Diamond p \rightarrow p}$
---	--

7. Axioms for \mathbb{U} and \mathbb{R} :

$$\begin{array}{ll} \text{(p)} \quad p \mathbb{U} q \rightarrow \diamond q & \text{(r)} \quad \Box q \rightarrow p \mathbb{R} q \\ \text{(q)} \quad p \mathbb{U} q \leftrightarrow (q \vee (p \wedge \circ (p \mathbb{U} q))) & \text{(s)} \quad p \mathbb{R} q \leftrightarrow (q \wedge (p \vee \widehat{\circ} (p \mathbb{R} q))) \end{array}$$

8. **Modus ponens:** $\frac{p \rightarrow q, p}{q}$

9. **Necessitation:** $\frac{p}{\widehat{\circ} p}$

Due to the intuitionistic basis of THT_ω all the axioms need to be duplicated in order to cope with the potential non-definability of modal operators (Simpson 1994). It should be also noticed that the induction axiom is replaced by the corresponding inference rule. These inference rules can be derived from the axioms and vice versa. The corresponding proof of soundness and completeness is achieved by adapting Goldblatt's proof for classical LTL_ω (Goldblatt 1992) to the here-and-there case: soundness is proved by checking that all axioms are valid in THT_ω and that the inference rules preserve validity. The proof of completeness relies on variations of well-known methods for proving completeness: a canonical model construction for bi-relational models (Simpson 1994), filtration and unwinding (Balbiani and Diéguez 2016).

From the axiomatic point of view, the main difference between THT_ω and THT_f relies on Axiom (3a), which guarantees that every time instant has a successor. From the point of view of the completeness proof, this axiom makes possible the unwinding of the filtered model by Balbiani and Diéguez (2016). Since Axiom (3a) is not valid in THT_f because it fails at the end of the trace, this unwinding may not be directly applied to obtain a sound and complete axiomatization of THT_f .

4 From TEL to automata

The relation between traditional temporal logics and automata is well-known and has been thoroughly studied in the literature (Demri et al. 2016). For infinite traces, LTL_ω unsatisfiability can be reduced to checking the emptiness of a language with a given automaton that accepts *infinite words*. For this, we can use, for instance, the well-known translation of Gerth et al. (1995) mapping a temporal formula φ into a *Büchi automaton* (Büchi 1962) that accepts the ω -regular language corresponding to the (infinite trace) models of φ . When we jump to finite traces, the type of automata required is simpler, as they just need to cover languages with words of finite length. Thus, in the works by Zhu et al. (2019) and De Giacomo and Vardi (2013), translations from LTL_f into different types of automata such as *Non-deterministic finite automata* (NFA) and *Alternating Automata on Words* or AFW (Chandra et al. 1981) have been proposed.

The use of automata for deciding THT -satisfiability can naturally follow the same steps, provided that the translation from Section 3.4 allows us to reduce the problem to standard LTL -satisfiability, regardless of the trace length. This fact was used by Cabalar and Diéguez (2014) to decide strong equivalence of two temporal formulas α and β in TEL_ω as follows. By Theorem 2, strong equivalence amounts to THT_ω -equivalence, that is, checking the validity of $\varphi := \alpha \leftrightarrow \beta$ in THT_ω . But then, by Theorem 4, this amounts to checking the unsatisfiability of $(21) \wedge \neg(\varphi^*)$ in LTL_ω and the latter is then done by the

standard Büchi-automata construction method. Similar steps can be followed for THT_f , replacing Büchi-automata for LTL_ω by NFA or AFW for LTL_f .

The use of automata for a *non-monotonic* temporal formalism like TEL becomes a more involved process, since temporal equilibrium models are obtained by a model selection (or minimization) process. The complexity of TEL_ω , for instance, suggests that there is no efficient method to reduce TEL_ω - to LTL_ω -satisfiability.

Theorem 5 (Bozzelli and Pearce 2015)

The problem of deciding whether a temporal formula has some (infinite) temporal equilibrium model in TEL_ω is EXPSpace-complete. □

A possible way of overcoming this difficulty is to observe that Definition 2 (of temporal equilibrium models) involves a kind of *quantification*. That is, we must find a total HT-model $\langle \mathbf{T}, \mathbf{T} \rangle$ of some formula such that *there is no smaller* $\mathbf{H} < \mathbf{T}$ such that $\langle \mathbf{H}, \mathbf{T} \rangle$ is also a model of the formula. This quantification seems to have a second-order nature, as it has to do with different configurations of the truth of propositions. In fact, in the non-temporal case, equilibrium models of a formula have been captured using Quantified Boolean Formulas (Pearce *et al.* 2001) and subsequently generalized by using a syntactic operator $\text{SM}[\cdot]$ (Ferraris *et al.* 2007). This gives a second-order formula that uses the translation from Section 3.4 and quantifies over primed propositions. This approach opens the natural possibility of capturing TEL through *Quantified LTL* (Sistla 1983), an extension of LTL in which we can use second-order quantifiers over propositions, and define a temporal version of the $\text{SM}[\cdot]$ operator as a QLTL formula. The advantage of having a quantified temporal formula $\text{SM}[\varphi]$ capturing the TS-models of φ is that there exist several methods for reducing that formula to an automaton, as we see below. Given that our interest is focused on finite automata, in the rest of this section, we assume that the alphabet \mathcal{A} is finite.

4.1 TS-models in terms of QLTL: the SM operator

We start by recalling the syntax and semantics of *Quantified Linear-Time Temporal Logic* or QLTL (Sistla 1983; Sistla *et al.* 1987; Demri *et al.* 2016). The syntax of QLTL extends the one of temporal formulas by quantifiers on propositional variables, that is, formulas of the type $\exists a \varphi$ and $\forall a \varphi$ are added.

Interpreting the additional formulas requires the concept of a variant trace.

Definition 6 (X-variant trace; French and Reynolds 2002)

Given two traces \mathbf{T} and \mathbf{T}' of length λ and $X \subseteq \mathcal{A}$, we say that \mathbf{T}' is an *X-variant trace* of \mathbf{T} if $T_i \setminus X = T'_i \setminus X$ for all $i \in [0.. \lambda)$.

With this, the semantics for QLTL is obtained from that of LTL by extending its satisfaction relation with the cases of quantified formulas as follows:

- $\mathbf{T}, i \models \forall a \varphi$ iff $\mathbf{T}', i \models \varphi$ for all $\{a\}$ -variant traces \mathbf{T}' of \mathbf{T} ,
- $\mathbf{T}, i \models \exists a \varphi$ iff $\mathbf{T}', i \models \varphi$ for some $\{a\}$ -variant traces \mathbf{T}' of \mathbf{T} .

It is known that QLTL_ω is as expressive as Büchi automata (Sistla *et al.* 1987), and so, more expressive than LTL_ω . For instance, the indicative “even states” property (Wolper

1983) can be expressed by the quantified temporal formula $\exists q (\neg q \wedge \Box (q \leftrightarrow \bigcirc \neg q)) \wedge \Box (q \rightarrow p)$.

We also recall that a quantified temporal formula is in *prenex normal form* if it is of the form $Q_1 x_1, Q_2 x_2 \cdots Q_n x_n \varphi$ where each $Q_i \in \{\exists, \forall\}$ and each x_i is a propositional variable occurring in the *quantifier-free* formula φ . An alternation in a quantified prefix is a sequence of $\exists y \forall x$ or $\forall x \exists y$ occurring in a prefix. Given a quantified temporal formula φ in prenex normal form, the *alternation depth* of φ is k if the prefix associated with φ contains k alternations. Also, given a quantified temporal formula φ , we denote by $free(\varphi)$ the number of propositional variables occurring in φ which are not quantified.

The complexity of the satisfiability problem in QLTL turns to be non-elementary (Sistla et al. 1987). However, for the case of formulas in prenex normal form with alternation depth k the following result has been obtained.

Theorem 6 (Sistla 1983)

The QLTL satisfiability problem of a quantified temporal formula φ in prenex normal form with alternation depth k is k -EXPSPACE-complete.⁷

Now, in order to encode TS-models of a temporal formula φ using a quantified temporal formula, we use translation φ^* from Section 3.4 and define the following notation. Let \mathbf{a} and \mathbf{a}' stand for the tuples $\langle a_1, \dots, a_n \rangle$ and $\langle a'_1, \dots, a'_n \rangle$ of n atoms, respectively. We define the following expressions:

$$\mathbf{a}' \leq \mathbf{a} \stackrel{def}{=} \bigwedge_{i=1}^n \Box (a'_i \rightarrow a_i) \qquad \mathbf{a}' < \mathbf{a} \stackrel{def}{=} \mathbf{a}' \leq \mathbf{a} \wedge \bigvee_{i=1}^n \Diamond (\neg a'_i \wedge a_i).$$

Note that, whenever \mathbf{a} coincides with the original propositional signature \mathcal{A} , $\mathbf{a}' \leq \mathbf{a}$ amounts to axiom (21) and is used to guarantee that atoms in trace \mathbf{H} are also included in trace \mathbf{T} . Expression $\mathbf{a}' < \mathbf{a}$ further forces $\mathbf{H} < \mathbf{T}$ as stated by the following proposition.

Proposition 15

Let \mathbf{T}^* be a trace over signature $\mathcal{A}^* = \mathcal{A} \cup \{a' \mid a \in \mathcal{A}\}$ and let \mathbf{H} and \mathbf{T} be such that $H_i \stackrel{def}{=} \{a \mid a' \in T_i^*\}$ and $T_i \stackrel{def}{=} T_i^* \cap \mathcal{A}$. Also, let \mathbf{a} be a tuple containing all propositional variables on \mathcal{A} and let \mathbf{a}' be of the same length as \mathbf{a} . Then, we have:

1. $\mathbf{T}^*, 0 \models \mathbf{a}' \leq \mathbf{a}$ iff $\mathbf{H} \leq \mathbf{T}$,
2. $\mathbf{T}^*, 0 \models \mathbf{a}' < \mathbf{a}$ iff $\mathbf{H} < \mathbf{T}$. \(\boxtimes\)

We can now define the temporal version of the SM operator by Ferraris et al. (2007): Let φ be a temporal formula over signature \mathcal{A} and let \mathbf{a} be a tuple with all propositions in \mathcal{A} . Then, we define the QLTL expression:

$$SM[\varphi] \stackrel{def}{=} \varphi \wedge \neg \exists \mathbf{a}' (\mathbf{a}' < \mathbf{a} \wedge \varphi^*). \tag{27}$$

The fact that $SM[\varphi]$ captures the TS-models of φ is established by the following theorem.

Theorem 7

A trace \mathbf{T} is a TS-model of a temporal formula φ iff \mathbf{T} is a QLTL-model of $SM[\varphi]$. \(\boxtimes\)

⁷ By convention, 0-EXPSPACE stands for PSPACE.

Note that a simple analysis of $SM[\varphi]$ allows us to match the upper bound complexity of the satisfiability problem of TEL obtained by Cabalar and Demri (2011).

Corollary 5

The satisfiability problem of TEL_ω is at most EXPSPACE-complete.

Proof

Note that $SM[\varphi]$ can be rewritten into $\forall \mathbf{a}' (\varphi \wedge (\mathbf{a}' < \mathbf{a} \rightarrow \neg\varphi^*))$ by applying the De Morgan laws and pushing φ , whose variables are free, inside of the scope of the second-order quantifiers. Since checking whether the previous formula is satisfiable iff $\exists \mathbf{a} \forall \mathbf{a}' (\varphi \wedge (\mathbf{a}' < \mathbf{a} \rightarrow \neg\varphi^*))$ is satisfiable (De Giacomo and Perelli 2021), we can apply Theorem 6 to obtain such EXPSPACE-complete upper bound. \square

Note that for finite traces this amounts to a characterization in terms of *weak QLTL*.

Example 3 (Example 2 continued)

Consider again formula $\varphi = \Box(\neg b \rightarrow a)$ in (23). As mentioned, formula $\mathbf{a} \leq \mathbf{a}'$ amounts to axiom (21) for all atoms in the signature which, in this case, corresponds to formula (22). Now $\mathbf{a} < \mathbf{a}'$ is stronger, as it requires not only $\mathbf{H} \leq \mathbf{T}$ but also $\mathbf{H} \neq \mathbf{T}$. In this case, we obtain the expression:

$$(22) \wedge (\Diamond(\neg a' \wedge a) \vee \Diamond(\neg b' \wedge b)), \tag{28}$$

whereas φ^* corresponds to (24), as illustrated above. So, as a result, $SM[\varphi]$ is finally unfolded into:

$$\begin{aligned} SM[\varphi] &= SM[\Box(\neg b \rightarrow a)] \\ &= \Box(\neg b \rightarrow a) \wedge \neg \exists a' b' (\\ &\quad \Box(a' \rightarrow a) \wedge \Box(b' \rightarrow b) \wedge (\Diamond(\neg a' \wedge a) \vee \Diamond(\neg b' \wedge b)) \\ &\quad \wedge \Box(\neg b \rightarrow a) \wedge \Box(\neg b \rightarrow a')). \end{aligned}$$

The LTL-models of $\Box(\neg b \rightarrow a)$ are traces \mathbf{T} where each state satisfies $T_i \neq \emptyset$. Take any \mathbf{T} where there is some T_i including atom b . Then, we can extend \mathbf{T} to \mathbf{T}^* for signature $\{a, b, a', b'\}$ repeating the truth of a' and b' with respect to a and b in all states, except in state T_i^* where we just leave b' false. It is not difficult to see that this extended interpretation \mathbf{T}^* satisfies $\mathbf{a} < \mathbf{a}' \wedge \varphi^*$ and so, \mathbf{T} cannot be a model of $SM[\varphi]$. Therefore, \mathbf{T} must make b false at all states and the only remaining model for $SM[\varphi]$ is the trace \mathbf{T} where $T_i = \{a\}$ for all states, which coincides with the only TS-model of φ .

4.2 From QLTL to automata

Once the TS-models of φ are captured by a quantified temporal formula, there are several possibilities to obtain the corresponding automaton. For instance, Cabalar and Demri (2011) used the following approach⁸ in the case of TEL_ω : First, we build a Büchi automaton \mathfrak{A}_1 that yields the LTL_ω -models of φ as usual. Then, we build a second automaton, \mathfrak{A}_2 , corresponding to the temporal formula $\mathbf{a} < \mathbf{a}' \wedge \varphi^*$ for the extended signature \mathcal{A}^* . On

⁸ In fact, Cabalar and Demri (2011) never constructed the formula $SM[\varphi]$ as an intermediate step, but the description we provide here is equivalent.

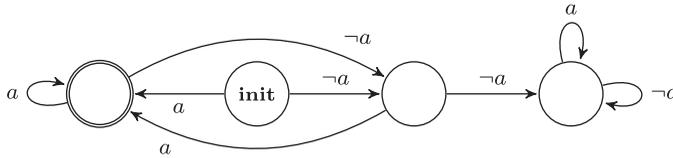


Fig. 1. NFA accepting the LTL_f-models of $\Box(\neg a \rightarrow \bigcirc a)$.

this last construction, we perform a filter operation to obtain the new automaton $h(\mathfrak{A}_2)$ that removes the primed atoms from the signature. In other words, $h(\mathfrak{A}_2)$ captures the models of the quantified temporal formula $\exists \mathbf{a}'(\mathbf{a} < \mathbf{a}' \wedge \varphi^*)$ and, in terms of THT, this corresponds to the **T** traces for which there exists some $\mathbf{H} < \mathbf{T}$. The final step is using the operations of complement and intersection of Büchi automata to obtain $\mathfrak{A}_1 \cap h(\mathfrak{A}_2)$ that naturally corresponds to the formula (27) of the SM operator. This method was implemented later on in the tool **abstem** (Cabalar and Diéguez 2014) that allows for both computing TS-models and checking strong equivalence of temporal formulas in TEL_ω .

In what follows, we present an analogous method for TEL_f yet by reducing SM formulas to automata over finite words. More precisely, we consider two different types of finite automata \mathfrak{A} , namely, NFAs and AFWs.

A NFA is a structure $(\Sigma, Q, Q_0, \delta, F)$ where

1. Σ is the alphabet,
2. Q is a set of states,
3. $Q_0 \subseteq Q$ is a set of initial states,
4. $\delta : Q \times \Sigma \mapsto 2^Q$ is a transition function and
5. $F \subseteq Q$ is a set of final states.

A run of a NFA for an input word $w = (X_0 \dots X_{n-1})$ of length $n > 0$ is a finite sequence of states q_0, \dots, q_n such that $q_0 \in Q_0$ and $\delta(q_i, X_i) = q_{i+1}$ for all $0 \leq i < n$. The run is said to be *accepting* if $q_n \in F$. By $\mathcal{L}(\mathfrak{A})$ we denote the set of runs accepted by \mathfrak{A} . The following theorem shows the relation between LTL_f and NFAs.

Theorem 8 (Zhu et al. 2017; Camacho et al. 2018)

A given temporal formula φ over \mathcal{A} can be translated into a NFA \mathfrak{A}_φ such that the set of LTL_f-models of φ corresponds to $\mathcal{L}(\mathfrak{A}_\varphi)$. □

Note that the alphabet of \mathfrak{A}_φ consists of $2^{\mathcal{A}}$, which amounts to the set of interpretations over \mathcal{A} .

To give an example of how the corresponding algorithm works, let us consider the formula $\varphi = \Box(\neg a \rightarrow \bigcirc a)$ in (2). In LTL_f, φ is equivalent to $\Box(a \vee \bigcirc a)$, which means that every LTL_f-model of this formula satisfies a at a state i or $i + 1$, for all $0 \leq i < \lambda$. The set of LTL_f-models of φ is captured by the NFA in Figure 1. We see that every run in which a is false in two consecutive states is disregarded.

On the other hand, an AFW is a structure $(\Sigma, Q, q_0, \delta, F)$ where:

1. Σ is the alphabet,
2. Q is a set of states,
3. $q_0 \in Q$ is the initial state,

4. $\delta : Q \times \Sigma \mapsto \mathbb{B}^+(Q)$ is a transition function, where $\mathbb{B}^+(Q)$ is built from the elements of Q , conjunction, disjunction, \top and \perp ,
5. $F \subseteq Q$ is a set of final states.

Given an input word $w = (X_0 \dots X_{n-1})$ of length $n > 0$, a run of an AFW is a tree labeled by states of the AFW such that

1. the root is labeled by q_0 ,
2. if a tree node z at level i is labeled by a state q and $\delta(q, X_i) = \Theta$ then either $\Theta = \top$ or some $P \subseteq Q$ satisfies Θ and z has a child for each element in P , and
3. the run is accepting if all leaves at depth n are labeled by states in F .

Thus, a branch in an accepting run has to hit the \top transition or hit an accepting state after reading all the input word w .

The correspondence between LTL_f and AFW is stated in the following theorem

Theorem 9 (De Giacomo and Vardi 2013)

A temporal formula φ can be translated into an AFW $\mathfrak{A}_\varphi = (2^{A \cup \{last\}}, Q, q_\varphi, \delta, \emptyset)$, where

1. *last* corresponds a fresh atom indicating the last state of the trace,
2. Q corresponds to the negation-closed closure of φ due to Fischer and Ladner (1979),
3. q_φ is the initial state, which corresponds to the state labeled with the formula φ ,
4. $\delta : Q \times 2^{A \cup \{last\}} \mapsto \mathbb{B}^+(Q)$ is the transition function defined as follows, where $X \subseteq A \cup \{last\}$ is a state

(a) $\delta(q_\top, X) = \top$

(b) $\delta(q_\perp, X) = \perp$

(c) $\delta(q_a, X) = \begin{cases} \top & \text{if } a \in X; \\ \perp & \text{otherwise} \end{cases}$

(d) $\delta(q_{\neg\varphi}, X) = \overline{\delta(q_\varphi, X)}$, where $\overline{\delta(q_\varphi, X)}$ is obtained from $\delta(q_\varphi, X)$ by switching \wedge and \vee , by switching \top and \perp and, in addition, by negating subformulas in X .

(e) $\delta(q_{\varphi \wedge \psi}, X) = \delta(q_\varphi, X) \wedge \delta(q_\psi, X)$

(f) $\delta(q_{\varphi \vee \psi}, X) = \delta(q_\varphi, X) \vee \delta(q_\psi, X)$

(g) $\delta(q_{\varphi \rightarrow \psi}, X) = \delta(q_{\neg\varphi}, X) \vee \delta(q_\psi, X)$

(h) $\delta(q_{\circ\varphi}, X) = \begin{cases} q_\varphi & \text{if } last \notin X \\ \perp & \text{otherwise} \end{cases}$

(i) $\delta(q_{\widehat{\circ}\varphi}, X) = \begin{cases} q_\varphi & \text{if } last \notin X \\ \top & \text{otherwise} \end{cases}$

(j) $\delta(q_{\varphi \cup \psi}, X) = \begin{cases} \delta(q_\psi, X) & \text{if } last \in X \\ \delta(q_\psi, X) \vee (\delta(q_\varphi, X) \wedge q_{\varphi \cup \psi}) & \text{otherwise} \end{cases}$

(k) $\delta(q_{\varphi \mathbb{R} \psi}, X) = \begin{cases} \delta(q_\psi, X) & \text{if } last \in X \\ \delta(q_\psi, X) \wedge (\delta(q_\varphi, X) \vee q_{\varphi \mathbb{R} \psi}) & \text{otherwise} \end{cases} \quad \square$

Moreover, as in the case of NFAs, $\mathcal{L}(\mathfrak{A}_\varphi)$ corresponds to the set of LTL_f -models of φ .

As above, the formula $\square(\neg a \rightarrow \circ a) \equiv \square(a \vee \circ a)$ can be translated into the AFW of Figure 2. Here, the alternating automaton is represented as a NFA extended with a new

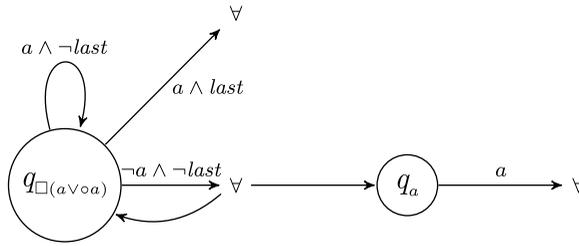


Fig. 2. AFW accepting the LTL_f -models of $\Box(a \vee \bigcirc a)$.

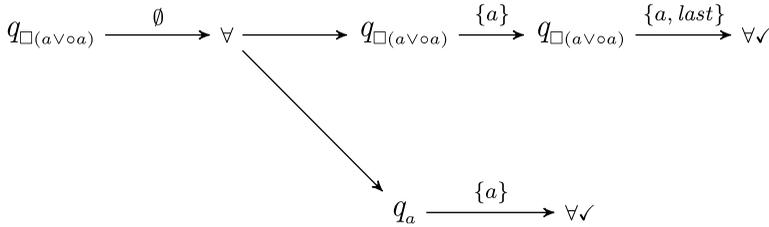


Fig. 3. A branching run of the AFW of Figure 2 accepting the input $\emptyset \cdot \{a\} \cdot \{a, last\}$.

node type, labeled as “ \forall ” and splitting some incoming edge into a (possibly empty) set of outgoing, unlabeled edges. As in the NFA, if multiple outgoing edges from a state share the same label X , they are understood as an *existential* constraint on paths (a run must follow at least one of them). As also happens in a NFA, when there is no outgoing edge for some label X , it is understood as a transition to \perp (the “empty” existential constraint), that is, no accepting path can be built using X . As it can be imagined, the new \forall -nodes represent the dual, *universal* constraints on paths: a run reaching \forall must follow *all* the outgoing edges. Analogously, when the \forall -node has no outgoing edges (the “empty” universal constraint \top) the run is trivially accepted.

Figure 3 shows an accepting run of the AFW of Figure 2 with the input $\emptyset \cdot \{a\} \cdot \{a, last\}$. For sake of clarity, we have labeled each tree edge with the input symbol X_i that fires the transition and, when applicable, we also show the \forall constraints that are reached along the run. As we can see, since $X_0 = \emptyset$ satisfies $\neg a \wedge \neg last$, it splits the run into two branches, but both of them eventually reach the empty \forall (i.e. the accepting \top transition).

To illustrate a non-accepting run, suppose we take the input $\emptyset \cdot \{a\} \cdot \{last\}$ instead, so that the last interpretation X_2 does not satisfy atom a now. Figure 4 shows the corresponding run obtained from the AFW in Figure 2. As before, $X_0 = \emptyset$ splits the run but, this time, the path followed by the top branch is not accepted since no outgoing edge from $q_{\Box(a \vee \bigcirc a)}$ has a formula satisfied by $X_2 = \{last\}$. This implicitly means $\delta(q_{\Box(a \vee \bigcirc a)}, \{last\}) = \perp$.

The rest of this section is devoted to show how Theorem 8 and 9 can be used to obtain the following result.

Theorem 10

A temporal formula φ can be translated into an AFW/NFA \mathfrak{A}_φ such that $\mathcal{L}(\mathfrak{A}_\varphi)$ corresponds to the set of TEL_f -models of φ .

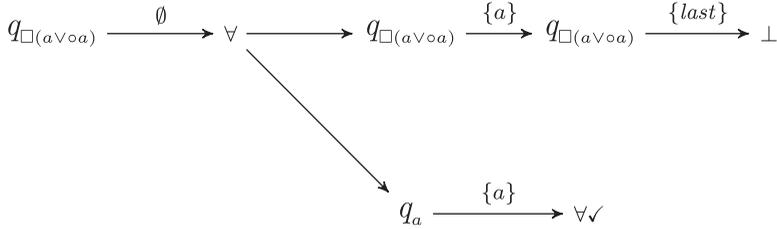


Fig. 4. A run of the AFW of Figure 2 rejecting the input $\emptyset \cdot \{a\} \cdot \{last\}$.

Our construction relies on several intermediate automata transformations corresponding to the ones by Cabalar and Demri (2011). Each of them accepts the model of one specific part of the SM expression. The involved automata and their correspondence with $SM[\varphi]$ are depicted in (29).

$$SM[\varphi] = \underbrace{\varphi}_{\mathfrak{A}_1} \wedge \neg \exists \mathbf{a}' \underbrace{\mathbf{a}' < \mathbf{a} \wedge (\varphi)^*}_{\mathfrak{A}_2}. \tag{29}$$

$$\underbrace{\hspace{10em}}_{h(\mathfrak{A}_2)}$$

$$\underbrace{\hspace{10em}}_{\overline{h(\mathfrak{A}_2)}}$$

$$\underbrace{\hspace{10em}}_{\mathfrak{A}_1 \cap \overline{h(\mathfrak{A}_2)}}$$

The computation of the final automaton $\mathfrak{A}_1 \cap \overline{h(\mathfrak{A}_2)}$ accepting the TEL_f -models of φ is done in a compositional manner, analogous to the translation of QLTL or *Weak Monadic Second Order Theory of one Successor* into NFA/AFW (Bozzelli et al. 2015; Henriksen et al. 1995). The two initial automata \mathfrak{A}_1 and \mathfrak{A}_2 are computed by using Lemma 9 or 8, depending on whether we are interested in computing an NFA or an AFW. Note that the formulas recognized by \mathfrak{A}_1 and \mathfrak{A}_2 are temporal formulas, which justifies the application of these lemmas. $h(\mathfrak{A}_2)$ is obtained from \mathfrak{A}_2 by *projecting* the atoms of the type a' . In language-theoretic terms, a *projection operation* $h(\mathcal{L}(\mathfrak{A}))$ is defined as follows.

$$h(\mathcal{L}(\mathfrak{A})) \stackrel{def}{=} \{w \mid \exists w' \text{ s.t. } w \text{ is identical to } w' \text{ except for all atoms of the type } a'\}.$$

Since $h(\mathcal{L}(\mathfrak{A})) = \mathcal{L}(h(\mathfrak{A}))$ (Cabalar and Demri 2011), the projection can be applied directly over \mathfrak{A}_2 and, broadly speaking, the resulting automaton accepts all LTL_f -models \mathbf{T} for which there exists a trace \mathbf{H} satisfying $\mathbf{H} < \mathbf{T}$. $\overline{h(\mathfrak{A}_2)}$ is obtained by complementing $h(\mathfrak{A}_2)$; it accepts all traces \mathbf{T} that are either no LTL_f -models of φ or for which exists no trace \mathbf{H} satisfying $\mathbf{H} < \mathbf{T}$. Automata complementation is the usual way to obtain automata accepting the LTL_f models of negated formulas. Finally, $\mathfrak{A}_1 \cap \overline{h(\mathfrak{A}_2)}$ is obtained by intersecting \mathfrak{A}_1 and $\overline{h(\mathfrak{A}_2)}$. This operation is equivalent to selecting all traces \mathbf{T} being LTL_f -models of φ for which there is no trace $\mathbf{H} < \mathbf{T}$ satisfying $\langle \mathbf{H}, \mathbf{T} \rangle, 0 \models \varphi$. In this case, $\langle \mathbf{T}, \mathbf{T} \rangle$ is an equilibrium model of φ .

If we analyze $\varphi = \Box(\neg a \rightarrow \circ a)$ in TEL_f , a is false at the initial state since it cannot be provable. Since $\neg a \rightarrow \circ a$ is satisfied in the initial state, a is true at time point 1, so the rule $\neg a \rightarrow \circ a$ cannot be applied to prove a at time point 2 and the cycle starts again. Hence, TEL_f -models of φ are of the form $\emptyset \cdot \{a\} \cdot \emptyset \cdot \{a\} \cdots$. They are captured by the NFA in Figure 5, whose accepting runs correspond to sequences of this type.

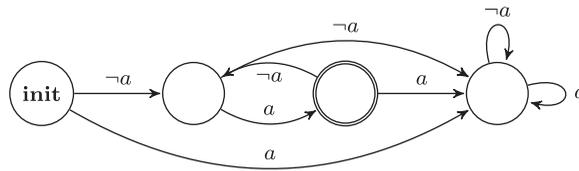


Fig. 5. NFA accepting the TEL_f -models of $\Box(\neg a \rightarrow \bigcirc a)$.

5 (Modal) temporal logic programs

In the previous section, we have seen a first method to compute TS-models and check TEL-satisfiability, relying first on a quantified LTL expression, and then, deriving the construction of different types of automata: Büchi for TEL_ω and NFA or AFW for TEL_f . Automata-based techniques are interesting for the analysis of *the set of traces* (TS-models) induced by a given temporal formula. This is useful, for instance, to check several types of equivalence between two different temporal representations, or to check the existence of a solution for a given planning problem. However, in many practical applications, rather than exploring all possibilities, our interest is more focused on the generation of one, or a few TS-models that encode solutions to a given temporal problem. This is, in fact, the usual model-based problem solving orientation of ASP, where stable models encode solutions to a particular problem, encoded by the rules in the logic program. In this way, if we want to solve a temporal diagnosis scenario, we are interested in finding at least one TS-model that explains the observations. Similarly, for a given planning problem, we look for some plan that reaches the goal in a finite number of steps.

In this section, we show how the generation of TS-models can be done using ASP technology by considering temporal theories with a syntax closer to logic programming. As we see next, both TEL_ω and TEL_f can be reduced to a normal form that we call *(modal) temporal logic programs*. This normal form is useful for ASP-based computation. For the finite case, there are two translations of this normal form into ASP. One allows any temporal formula and provides a translation for a given trace length. The other translation imposes some restrictions on the formula but allows for incremental solving. There are tools supporting these translations, most notably `telingo` (Cabalar et al. 2019) for the incremental approach.

We provide a generic normal form for temporal programs and a translation for temporal formulas into this normal form. Both the finite and the infinite case are special cases of this normal form.

Definition 7 (Temporal literal, rule, and program)

Given alphabet \mathcal{A} , we define the set of *temporal literals* as $\{a, \neg a, \bullet a, \neg \bullet a \mid a \in \mathcal{A}\}$.

A *temporal rule* is either:

- an *initial rule* of the form $B \rightarrow A$
- a *dynamic rule* of the form $\widehat{\Box}(B \rightarrow A)$
- a *fulfillment rule* of the form $\Box(\Box p \rightarrow q)$ or $\Box(p \rightarrow \Diamond q)$
- a *final rule* of the form $\Box(\mathbb{F} \rightarrow (B \rightarrow A))$

where $B = b_1 \wedge \dots \wedge b_n$ with $n \geq 0$, $A = a_1 \vee \dots \vee a_m$ with $m \geq 0$ and the b_i and a_j are temporal literals for dynamic rules and regular literals $\{a, \neg a \mid a \in \mathcal{A}\}$ for initial and

Table 1. Definition of formulas

μ	$df(\mu)$
$\varphi \wedge \psi$	$\Box(\ell_\mu \leftrightarrow \ell_\varphi \wedge \ell_\psi)$
$\varphi \vee \psi$	$\Box(\ell_\mu \leftrightarrow \ell_\varphi \vee \ell_\psi)$
$\varphi \rightarrow \psi$	$\Box(\ell_\mu \leftrightarrow \ell_\varphi \rightarrow \ell_\psi)$
$\circ\varphi$	$\widehat{\Box}(\bullet\ell_\mu \leftrightarrow \ell_\varphi)$ $\Box(\Box\ell_\mu \rightarrow \ell_{\neg\Diamond\mathbb{F}})$
$\bullet\varphi$	$\neg\ell_\mu$ $\widehat{\Box}(\ell_\mu \leftrightarrow \bullet\ell_\varphi)$

Table 2. Definition of formulas

μ	$df(\mu)$
$\varphi \cup \psi$	$\widehat{\Box}(\bullet\ell_\mu \leftrightarrow \bullet\ell_\psi \vee (\bullet\ell_\varphi \wedge \ell_\mu))$ $\Box(\ell_\mu \rightarrow \Diamond\ell_\psi)$ $\Box(\ell_\psi \rightarrow \Diamond\ell_\mu)$
$\varphi \mathbb{R} \psi$	$\widehat{\Box}(\bullet\ell_\mu \leftrightarrow \bullet\ell_\psi \wedge (\bullet\ell_\varphi \vee \ell_\mu))$ $\Box(\Box\ell_\psi \rightarrow \ell_\mu)$ $\Box(\Box\ell_\mu \rightarrow \ell_\psi)$
$\varphi \mathbf{S} \psi$	$\ell_\mu \leftrightarrow \ell_\psi$ $\widehat{\Box}(\ell_\mu \leftrightarrow \ell_\psi \vee (\ell_\varphi \wedge \bullet\ell_\mu))$
$\varphi \mathbf{T} \psi$	$\ell_\mu \leftrightarrow \ell_\psi$ $\widehat{\Box}(\ell_\mu \leftrightarrow \ell_\psi \wedge (\ell_\varphi \vee \bullet\ell_\mu))$

final rules, and p and q are atoms. We call *temporal logic program* to a set of temporal rules.

We normally allow some straightforward extensions of this form. For instance, an *always-rule* has the form $\Box(B \rightarrow A)$ and stands for the conjunction $(B \rightarrow A) \wedge \widehat{\Box}(B \rightarrow A)$ of an initial and a dynamic rule, respectively (therefore, B and A only contain regular literals). Moreover, an *extended dynamic rule* allows body literals b_i to be the atomic formulas \mathbf{I} or \mathbb{F} (since they can be encoded as auxiliary atoms). When we allow this last extension, final rules can be reduced to constraints (empty head) of the form $\Box(\mathbb{F} \rightarrow (B \rightarrow \perp))$.

Theorem 11 (Normal form; Aguado et al. 2013; Cabalar et al. 2018)

Every temporal formula φ can be converted into a temporal program being THT_f -equivalent to φ and into one being THT_ω -equivalent to φ , in both cases, modulo auxiliary atoms.

Moreover, the reduction by Cabalar et al. (2018) further guarantees the following property.

Corollary 6

For finite traces $\lambda \in \mathbb{N}$, every temporal formula can be converted into a THT_f -equivalent temporal logic program P consisting of initial, dynamic, and final rules only.

We denote these three sets of rules as $I(P)$, $D(P)$ and $F(P)$, respectively. □

In what follows, we introduce a more general reduction from temporal formulas to temporal logic programs that is applicable to *both* finite and infinite traces. It uses an extended alphabet $\mathcal{A}^+ \supseteq \mathcal{A}$ that additionally contains a new atom ℓ_φ (a.k.a. label) for each formula φ in the original language over \mathcal{A} along with the label $\ell_{\neg\Diamond\mathbb{F}}$. This label represents the value of the formula $\neg\Diamond\mathbb{F}$, which means that it can be used to restrict ourselves to finite traces by including the formula $\Box(\neg\ell_{\neg\Diamond\mathbb{F}})$. For convenience, we use $\ell_\varphi \stackrel{def}{=} \varphi$ if φ is \top, \perp or an atom $a \in \mathcal{A}$. For any non-atomic formula μ over \mathcal{A} , we define the translation df given in Tables 1 and 2, and call df the *definition* of μ . Note that the translation of the *next* operator \circ includes the label $\ell_{\neg\Diamond\mathbb{F}}$. This ensures that if the formula $\Box(\neg\ell_{\neg\Diamond\mathbb{F}})$ is included and therefore only finite traces are considered, then $\circ\varphi$ cannot be satisfied in the final state.

Table 3. Translation of Γ_1 into normal form

μ	$df^*(\mu)$	μ	$df^*(\mu)$
$\neg p$	$\ell_1 \wedge p \rightarrow \perp$ (30)	$\neg p \rightarrow q \cup p$	$\ell_3 \wedge \ell_1 \rightarrow \ell_2$ (39)
	$\neg p \rightarrow \ell_1$ (31)		$\neg \ell_1 \rightarrow \ell_3$ (40)
	$\widehat{\square}(\ell_1 \wedge p \rightarrow \perp)$ (32)		$\ell_2 \rightarrow \ell_3$ (41)
	$\widehat{\square}(\neg p \rightarrow \ell_1)$ (33)		$\ell_1 \vee \neg \ell_2 \vee \ell_3$ (42)
$q \cup p$	$\widehat{\square}(\bullet \ell_2 \rightarrow \bullet p \vee \bullet q \wedge \ell_2)$ (34)		$\widehat{\square}(\ell_2 \rightarrow \ell_3)$ (43)
	$\widehat{\square}(\bullet q \wedge \ell_2 \rightarrow \bullet \ell_2)$ (35)		$\widehat{\square}(\ell_3 \wedge \ell_1 \rightarrow \ell_2)$ (44)
	$\widehat{\square}(\bullet p \rightarrow \bullet \ell_2)$ (36)	$\widehat{\square}(\neg \ell_1 \rightarrow \ell_3)$ (45)	
	$\square(\ell_2 \rightarrow \diamond p)$ (37)	$\widehat{\square}(\top \rightarrow \ell_1 \vee \neg \ell_2 \vee \ell_3)$ (46)	
	$\square(p \rightarrow \diamond \ell_2)$ (38)	$\square(\neg p \rightarrow q \cup p)$	$\widehat{\square}(\bullet \ell_3 \wedge \ell_4 \rightarrow \bullet \ell_4)$ (47)
			$\widehat{\square}(\bullet \ell_4 \rightarrow \bullet \ell_3)$ (48)
	$\widehat{\square}(\bullet \ell_4 \rightarrow \ell_4)$ (49)		
	$\square(\square \ell_3 \rightarrow \ell_4)$ (50)		
	$\square(\square \ell_4 \rightarrow \ell_3)$ (51)		

These definitions of temporal formulas are not yet in normal form, but they contain some double implications that can easily be transformed into the format of temporal logic programs by simple, non-modal transformations in the propositional logic of HT. This translation of the definition into normal form was presented by Cabalar et al. (2018) for the finite case and by Aguado et al. (2013) for the infinite case. We denote the resulting set of formulas in normal form by df^* . Given a theory Γ , we define $sub(\Gamma)$ as the set of all subformulas of all formulas in Γ .

Definition 8 (Aguado et al. 2013; Cabalar et al. 2018)

We define the translation σ as the following temporal logic program:

$$\sigma(\Gamma) = \{\ell_\gamma \mid \gamma \in \Gamma\} \cup \{df^*(\mu) \mid \mu \in sub(\Gamma)\}. \quad \boxtimes$$

Theorem 12 (Aguado et al. 2013; Cabalar et al. 2018)

For any theory Γ over \mathcal{A} , we have $\{\mathbf{M} \mid \mathbf{M} \models \Gamma\} = \{\mathbf{M}' \mid_{\mathcal{A}} \mathbf{M}' \models \sigma(\Gamma)\}. \quad \boxtimes$

This correspondence between models of Γ and $\sigma(\Gamma)$ is, in fact, one-to-one, since the satisfaction of each label ℓ_γ is equivalent to the satisfaction of its associated formula γ . The next result shows that the computation of $\sigma(\Gamma)$ has a polynomial complexity on the size of Γ .

Theorem 13

Translation σ is linear. \boxtimes

As an example, consider the theory $\Gamma_1 = \square(\neg p \rightarrow q \cup p)$. Its translation $\sigma(\Gamma_1)$ can be determined using Table 3, independently of the finiteness of the trace. On finite traces, the fulfillment rules (37), (38), (50) and (51) can be replaced by the corresponding final rules. On infinite traces, the initial rules (30), (31), (39), (40), (41), (42) can be combined with dynamic rules (32), (33), (43), (44), (45), (46) into formulas of the form $\square(r)$ where r is the given initial rule. Rules (38) and (51) are superfluous in the infinite case.

The interest of temporal logic programs is that, as suggested by their name, they have a syntax closer to logic programming. This suggests that, at least for finite traces as in TEL_f , computing TS-models can be delegated to an ASP solver in some way. In the rest of the section, we explain two different methods to achieve this goal.

5.1 Bounded translation of temporal programs over finite traces to ASP

A first way to encode a temporal program P into standard ASP is by assuming that we know the (finite) trace length λ beforehand. In that case, we can compute all models in $\text{TEL}(P, \lambda)$ by a simple translation of P into a regular program. For this, we let $\mathcal{A}_k = \{a_k \mid a \in \mathcal{A}\}$ be a time stamped copy of alphabet \mathcal{A} for each time point $k \in [0..\lambda)$.

Definition 9 (Bounded translation; Cabalar et al. 2018)

We define the translation τ of a temporal literal at time point k as:

$$\begin{aligned} \tau_k(a) &\stackrel{\text{def}}{=} a_k & \tau_k(\neg a) &\stackrel{\text{def}}{=} \neg a_k & \text{for } a \in \mathcal{A} \\ \tau_k(\bullet a) &\stackrel{\text{def}}{=} a_{k-1} & \tau_k(\neg \bullet a) &\stackrel{\text{def}}{=} \neg a_{k-1} & \text{for } a \in \mathcal{A}. \end{aligned}$$

We define the translation at time point k of any (non-fulfillment) rule r as those in Definition 7 as:

$$\tau_k(r) \stackrel{\text{def}}{=} \tau_k(a_1) \vee \dots \vee \tau_k(a_m) \leftarrow \tau_k(b_1) \wedge \dots \wedge \tau_k(b_n).$$

We define the translation of a temporal program P bounded by finite length λ as:

$$\tau_\lambda(P) \stackrel{\text{def}}{=} \{\tau_0(r) \mid r \in I(P)\} \cup \{\tau_k(r) \mid r \in D(P), k \in [1..\lambda)\} \cup \{\tau_{\lambda-1}(r) \mid r \in F(P)\},$$

where $I(P), D(P)$ and $F(P)$ are the initial, dynamic and final rules in P . □

Note that the translation of temporal rules is similar in just considering the implication $B \rightarrow A$ in Definition 7; their difference manifests itself in their instantiation in $\tau_\lambda(P)$.

As an example, let program P be the set of temporal rules:

$$\{ \rightarrow a, \quad \widehat{\circ} \square(\bullet a \rightarrow b), \quad \square(\mathbb{F} \rightarrow (\neg b \rightarrow \perp)) \}. \tag{52}$$

This program has a single finite temporal stable model of length 2, viz. $\{a\} \cdot \{b\}$. Applying translation τ for some bound λ to our temporal program P in (52) yields regular logic programs of the following form.

$$\tau_\lambda(P) = \{a_0 \leftarrow\} \cup \{b_k \leftarrow a_{k-1} \mid k \in [1..\lambda)\} \cup \{\perp \leftarrow \neg b_{\lambda-1}\}.$$

Program $\tau_1(P)$ has the stable model $\{a_0, b_1\}$ but all $\tau_\lambda(P)$ for $\lambda > 2$ are unsatisfiable (have no stable models).

Theorem 14 (Cabalar et al. 2018)

Let P be a temporal program over \mathcal{A} . Let $\mathbf{T} = (T_i)_{i \in [0..\lambda)}$ be a trace of finite length λ over \mathcal{A} and X a set of atoms over $\bigcup_{i \in [0..\lambda)} \mathcal{A}_i$ such that $a \in T_i$ iff $a_i \in X$ for $i \in [0..\lambda)$.

Then, \mathbf{T} is a temporal stable model of P iff X is a stable model of $\tau_\lambda(P)$. □

Applied to our example, this result confirms that the temporal stable model $\{a\} \cdot \{b\}$ of P corresponds to the stable model $\{a_0, b_1\}$ of $\tau_2(P)$.

Using this translation we have implemented a system, **tel**⁹, that takes a propositional theory Γ of arbitrary temporal formulas and a bound λ and returns the regular logic program $\tau_\lambda(P)$, where P is the intermediate normal form of Γ left implicit. The resulting program $\tau_\lambda(P)$ can then be solved by any off-the-shelf ASP system. For illustration, consider the representation of our example temporal program in (52) in **tel**'s input language.

⁹ <https://github.com/potassco/tel>

```

a.
#next^ #always+ ( (#previous a) -> b).
#always+ ( #final -> (~ b -> #false)).

```

As expected, passing the result of `tel`'s translation for horizon 2 to `clingo` yields the stable model containing `a(0)` and `b(1)` (suppressing auxiliary atoms).

5.2 Pointwise translation of temporal programs over finite traces to ASP

The bounded translation $\tau_\lambda(P)$ allows us to compute all models in $\text{TEL}(P, \lambda)$ for a fixed bound λ . However, in many practical problems (as in planning, for instance), λ is unknown beforehand and the crucial task consists in finding a representation of $\text{TEL}(P, k)$ that is easily obtained from that of $\text{TEL}(P, k-1)$. In ASP, this can be accomplished via incremental solving techniques that rely upon the composition of logic program modules (Oikarinen and Janhunen 2006). The idea is then to associate the knowledge at each time point with a module and to successively add modules corresponding to increasing time points (while leaving all previous modules unchanged). A stable model obtained after k compositions then corresponds to a TEL_f -model of length k . This technique of modular computation, however, is only applicable when modules are *compositional* (positive loops cannot be formed across modules), something that cannot always be guaranteed for arbitrary temporal programs. Still, Cabalar et al. (2018) identified a quite general syntactic fragment¹⁰ that implies compositionality. We say that a temporal rule as in Definition 7 is *present-centered*, whenever all the literals a_1, \dots, a_m in its head A are regular. Accordingly, a set of such rules is a present-centered temporal program. In fact, such programs are sufficient to capture common action languages. For illustration, we show how to encode action language \mathcal{BC} (Lee et al. 2013) into present-centered temporal programs in TEL_f in Section 6.

Following these ideas, we provide next a “point-wise” variant of our translation that allows for defining one module per time point and is compositional for the case of present-centered temporal programs. We begin with some definitions. A *module* \mathbb{P} is a triple (P, I, O) consisting of a logic program P over alphabet \mathcal{A}_P and sets I and O of *input* and *output* atoms such that

1. $I \cap O = \emptyset$,
2. $\mathcal{A}_P \subseteq I \cup O$, and
3. $H(P) \subseteq O$,

where $H(P)$ gives all atoms occurring in rule heads in P . Whenever clear from context, we associate \mathbb{P} with (P, I, O) . In our setting, a set X of atoms is a stable model of \mathbb{P} , if X is a stable model of logic program P .¹¹ Two modules \mathbb{P}_1 and \mathbb{P}_2 are *compositional*, if $O_1 \cap O_2 = \emptyset$ and $O_1 \cap C = \emptyset$ or $O_2 \cap C = \emptyset$ for every strongly connected component C of the positive dependency graph of the logic program $P_1 \cup P_2$. In other words, all rules defining an atom must belong to the same module, and no positive recursion is allowed among modules. Whenever \mathbb{P}_1 and \mathbb{P}_2 are compositional, their *join* is defined as the module

¹⁰ In order to compute loop formulas for TEL_ω , Cabalar and Diéguez (2011) used a similar fragment (*splittable programs*) where rules cannot derive information from the future to the past.

¹¹ Note that the default value assigned to input atoms is *false* in multi-shot solving (Gebser et al. 2019); this differs from the original definition (Oikarinen and Janhunen 2006) where a choice rule is used.

$\mathbb{P}_1 \sqcup \mathbb{P}_2 = (P_1 \cup P_2, (I_1 \setminus O_2) \cup (I_2 \setminus O_1), O_1 \cup O_2)$. The module theorem (Oikarinen and Janhunen 2006) ensures that compatible stable models of \mathbb{P}_1 and \mathbb{P}_2 can be combined to one of $\mathbb{P}_1 \sqcup \mathbb{P}_2$, and vice versa.

For literals and rules, the point-wise translation τ^* coincides with τ up to final rules.

Definition 10 (Point-wise translation: Temporal rules; Cabalar et al. 2018)

We define the translation of a final rule r as in Definition 7 at time point k as

$$\tau_k^*(r) \stackrel{\text{def}}{=} \tau_k(a_1) \vee \dots \vee \tau_k(a_m) \leftarrow \tau_k(b_1) \wedge \dots \wedge \tau_k(b_n) \wedge \neg q_{k+1}, \tag{53}$$

for a new atom $q \notin \mathcal{A}$ and of an initial or dynamic rule r as $\tau_k^*(r) \stackrel{\text{def}}{=} \tau_k(r)$.

The new atoms q_{k+1} in (53) are used to deactivate instances of final rules. This allows us to implement operator \mathbb{F} by using $\neg q_{k+1}$ and therefore to enable the actual final rule unless q_{k+1} is derivable. The idea is then to make sure that at each horizon k the atom q_{k+1} is false, while q_1, \dots, q_k are true. In this way, only $\tau_k^*(r)$ is potentially applicable, while all rules $\tau_i^*(r)$ are disabled at earlier time points $i \in [1..k]$.

Translation τ^* is then used to define modules for each time point as follows.

Definition 11 (Point-wise translation: Modules)

Let P be a present-centered temporal program over \mathcal{A} . We define the module \mathbb{P}_k corresponding to P at time point k as:

$$\mathbb{P}_0 \stackrel{\text{def}}{=} (P_0, \{q_1\}, \mathcal{A}_0) \quad \mathbb{P}_k \stackrel{\text{def}}{=} (P_k, \mathcal{A}_{k-1} \cup \{q_{k+1}\}, \mathcal{A}_k \cup \{q_k\}) \quad \text{for } k > 0,$$

where

$$\begin{aligned} P_0 &\stackrel{\text{def}}{=} \{\tau_0^*(r) \mid r \in I(P)\} \cup \{\tau_0^*(r) \mid r \in F(P)\} \\ P_k &\stackrel{\text{def}}{=} \{\tau_k^*(r) \mid r \in D(P)\} \cup \{\tau_k^*(r) \mid r \in F(P)\} \cup \{q_k \leftarrow\}. \end{aligned}$$

□

Each module \mathbb{P}_k defines what holds at time point k . The underlying present-centeredness warrants that modules only incorporate atoms from previous time points, as reflected by \mathcal{A}_{k-1} in the input of \mathbb{P}_k . The exception consists of auxiliary atoms like q_{k+1} that belong to the input of each \mathbb{P}_k for $k > 0$ but only get defined in the next module \mathbb{P}_{k+1} . This corresponds to the aforementioned idea that q_{k+1} is false when \mathbb{P}_k is the final module, and is set permanently to true once the horizon is incremented by adding \mathbb{P}_{k+1} . Note that atoms like q_{k+1} only occur negatively in rule bodies in \mathbb{P}_k and hence cannot invalidate the modularity condition. This technique allows us to capture the transience of final rules.

The point-wise translation of our present-centered example program P from (52) yields the following modules.

$$\begin{aligned} \mathbb{P}_0 &= (\{a_0 \leftarrow\} \cup \{\leftarrow \neg b_0, \neg q_1\}, \{q_1\}, \{a_0, b_0\}) \\ \mathbb{P}_i &= (\{b_i \leftarrow a_{i-1}\} \cup \{\leftarrow \neg b_i, \neg q_{i+1}\} \cup \{q_i \leftarrow\}, \{a_{i-1}, b_{i-1}, q_{i+1}\}, \{a_i, b_i, q_i\}) \\ \bigsqcup_{i=0}^{\lambda-1} \mathbb{P}_i &= (P_0 \cup \bigcup_{i=1}^{\lambda-1} P_i, \{q_\lambda\}, \{a_i, b_i \mid i \in [0..\lambda]\} \cup \{q_i \mid i \in [1..\lambda]\}). \end{aligned}$$

As above, only the composed module for $\lambda = 1$ yields a stable model, viz. $\{a_0, b_1, q_1\}$.

Theorem 15

Let P be a present-centered temporal program over \mathcal{A} . Let $\mathbf{T} = (T_i)_{i \in [0..\lambda]}$ be a trace of finite length λ over \mathcal{A} and X a set of atoms over $\bigcup_{0 \leq i < \lambda} \mathcal{A}_i$ such that $a \in T_i$ iff $a_i \in X$ for $i \in [0..\lambda]$. Then, we have that

\mathbf{T} is a temporal stable model of P iff $X \cup \{q_i \mid i \in [0..\lambda]\}$ is a stable model of $\bigsqcup_{i \in [0..\lambda]} \mathbb{P}_i$.

As with Theorem 14, this result confirms that the temporal stable model $\{a\} \cdot \{b\}$ of P corresponds to the stable model $\{a_0, b_1, q_1\}$ of $\mathbb{P}_0 \sqcup \mathbb{P}_1$.

As one might expect, not any temporal theory is reducible to a present-centered temporal program. Hence, computing models via incremental solving imposes some limitations on the possible combinations of temporal operators. Fortunately, Cabalar et al. (2018) identified again a quite natural and expressive syntactic fragment that is always reducible to present-centered programs. We say that a temporal formula is a *past–future rule* if it consists of rules as those in Definition 7 where B and A are just temporal formulas with the following restrictions: B and A contain no implications other than negations ($\alpha \rightarrow \perp$), B contains no future operators, and A contains no past operators. An example of a past–future rule is, for instance,

$$\Box(\text{shoot} \wedge \blacklozenge \text{shoot} \wedge \blacksquare \text{unloaded} \rightarrow \diamond \text{fail}), \quad (54)$$

capturing the sentence: “If we make two shots with a gun that was never loaded, then it will eventually fail.” Then, we have the following result.

Theorem 16 (Past–future reduction)

Every past–future rule r can be converted into a present-centered temporal program that is TEL_f -equivalent to r . \square

This past–future form is aligned with the orientation by Gabbay (1987a) where past operators in the body (or antecedent) are used to declaratively check the recorded story, whereas future operators in the head (or consequent) are employed to provide imperative commands to be executed afterwards, as a result of firing the rule.

We have implemented a second system, `telingo`¹² (Cabalar et al. 2019), that deals with present-centered temporal programs that are expressible in the full (non-ground) input language of `clingo` extended with temporal operators. In addition, `telingo` offers several syntactic extensions to facilitate temporal modeling: first, next operators can be used in singular heads and, second, arbitrary temporal formulas can be used in integrity constraints. All syntactic extensions beyond the normal form of Theorem 11 are compiled away by means of the translation used in its proof. The resulting present-centered temporal programs are then processed according the point-wise translation.

To facilitate the use of operators \bullet and \circ , `telingo` allows us to express them by adding leading or trailing quotes to the predicate names of atoms, respectively. For instance, the temporal literals $\bullet p(a)$ and $\circ q(b)$ can be expressed by `'p(a)` and `q'(b)`, respectively. For another example, consider the representation of the sentence “A robot cannot lift a box unless its capacity exceeds the box’s weight plus that of all held objects”:

```
:- lift(R,B), robot(R), box(B,W),
   #sum { C : capacity(R,C); -V,0 : 'holding(R,0,V) } < W.
```

Atom `'holding(R,0,V)` expresses what the robot was holding at the *previous* time point.

The distinction between different types of temporal rules is done in `telingo` via `clingo`’s `#program` directives (Gebser et al. 2019), which allow us to partition programs into subprograms. More precisely, each rule in `telingo`’s input language is associated

¹² <https://github.com/potassco/telingo>

Table 4. Two alternative *telingo* encodings for the temporal program in (52)

<pre>#program initial. a. #program dynamic. b :- 'a. #program final. :- not b.</pre>	<pre>#program always. a :- &initial. b :- 'a. :- not b, &final.</pre>
--	---

with a temporal rule r of form $(b_1 \wedge \dots \wedge b_n \rightarrow a_1 \vee \dots \vee a_m)$ as in Definition 7 and interpreted as $r, \widehat{\square}r$, or $\square(\mathbb{F} \rightarrow r)$ depending on whether it occurs in the scope of a program declaration headed by `initial`, `dynamic`, or `final`, respectively. Additionally, *telingo* offers `always` for gathering rules preceded by \square (thus dropping $\widehat{\square}$ from dynamic rules). A rule outside any such declaration is regarded to be in the scope of `initial`. This allows us to represent the temporal program in (52) in the two alternative ways shown in Table 4.

As mentioned, *telingo* allows us to use nested temporal formulas in integrity constraints as well as in negated form in place of temporal literals within rules. This is accomplished by encapsulating temporal formulas like φ in expressions of the form `&tel { φ }`. To this end, the full spectrum of temporal operators is at our disposal. They are expressed by operators built from `<` and `>` depending on whether they refer to the past or the future, respectively. So, `</1`, `<?/2`, and `<*/2` stand for \bullet , \mathbf{S} , and \mathbf{T} , and `>/1`, `>?/2`, `>*/2` for \circ , \cup , \mathbb{R} . Accordingly, `<*/1`, `<?/1`, `<:/1` represent \blacksquare , \blacklozenge , $\widehat{\circ}$, and analogously their future counterparts. \mathbf{I} and \mathbb{F} are represented by `&initial` and `&final`. This is complemented by Boolean connectives `&`, `|`, `~`, etc. For example, the integrity constraint ¹³ `'shoot \wedge \blacksquare unloaded \wedge \blacklozenge shoot \rightarrow \perp '` can be expressed as follows.

```
:- shoot, &tel { <* unloaded & < <? shoot }.
```

Once *telingo* has translated an (extended) temporal program into a regular one, it is incrementally solved by *clingo*'s multi-shot solving engine (Gebser *et al.* 2019).

To conclude this section, we provide a larger example of *telingo* encoding (a more detailed description of the system was presented by Cabalar *et al.* (2019)). Listing 1 contains an exemplary *telingo* encoding of the *Fox, Goose and Beans Puzzle* available at <https://github.com/potassco/telingo/tree/master/examples/river-crossing>.

Once upon a time a farmer went to a market and purchased a fox, a goose, and a bag of beans. On his way home, the farmer came to the bank of a river and rented a boat. But crossing the river by boat, the farmer could carry only himself and a single one of his purchases: the fox, the goose, or the bag of beans. If left unattended together, the fox would eat the goose, or the goose would eat the beans. The farmer's challenge was to carry himself and his purchases to the far bank of the river, leaving each purchase intact. How did he do it?

(https://en.wikipedia.org/wiki/Fox,_goose_and_bag_of_beans_puzzle)

In Listing 1, lines 3–5 and 9–10 provide facts holding in all and the initial states, respectively; this is indicated by the program directives headed by `always` and `initial`. The dynamic rules in lines 14–22 describe the transition function. The `farmer` moves

¹³ Similarly, formula (54) could be represented by several present-centered rules (including auxiliary atoms).

```

1 #program always.

3 item(fox;beans;goose).
4 route(river_bank,far_bank). route(far_bank,river_bank).
5 eats(fox,goose). eats(goose,beans).

7 #program initial.

9 at(farmer,river_bank).
10 at(X,river_bank) :- item(X).

12 #program dynamic.

14 move(farmer).
15 0 { move(X) : item(X) } 1.

17 at(X,B) :- 'at(X,A), move(X), route(A,B).
18 :- move(X), item(X), 'at(farmer,A), not 'at(X,A).

20 at(X,A) :- 'at(X,A), not move(X).

22 #program always.

24 :- at(X,A), at(X,B), A<B.
25 :- eats(X,Y), at(X,A), at(Y,A), not at(farmer,A).

27 #program final.

29 :- at(X,river_bank).

31 #show move/1.
32 #show at/2.

```

Listing 1. *telingo* encoding for the Fox, Goose and Beans Puzzle

at each time step (Line 14), and may take an item or not (Line 15). Line 17 describes the effect of action `move/1`, Line 18 its precondition, and Line 20 the law of inertia. The second part of the `always` rules give state constraints in Line 24 and 25. The `final` rule in Line 29 gives the goal condition. All in all, we obtain two shortest plans consisting of eight states in about 20 ms. Restricted to the move predicate, *telingo* reports the following solutions:

Time	Solution 1		Solution 2	
1				
2	<code>move(farmer)</code>	<code>move(goose)</code>	<code>move(farmer)</code>	<code>move(goose)</code>
3	<code>move(farmer)</code>		<code>move(farmer)</code>	
4	<code>move(beans)</code>	<code>move(farmer)</code>	<code>move(farmer)</code>	<code>move(fox)</code>
5	<code>move(farmer)</code>	<code>move(goose)</code>	<code>move(farmer)</code>	<code>move(goose)</code>
6	<code>move(farmer)</code>	<code>move(fox)</code>	<code>move(beans)</code>	<code>move(farmer)</code>
7	<code>move(farmer)</code>		<code>move(farmer)</code>	
8	<code>move(farmer)</code>	<code>move(goose)</code>	<code>move(farmer)</code>	<code>move(goose)</code>

This example was also used by Cabalar and Diéguez (2011) to illustrate the working of `stelp`, a tool for temporal ASP with TEL_ω . We note that `stelp` and `telingo` differ syntactically in describing transitions by using next or previous operators, respectively. Since `telingo` extends `clingo`'s input language, it offers a richer input language, as witnessed by the cardinality constraints in Line 15 in Listing 1. Finally, `stelp` uses a model checker and outputs an automaton capturing all infinite traces while `telingo` returns finite traces corresponding to plans.

6 TEL_f and the action language \mathcal{BC}

Present-centered temporal programs over finite traces are sufficient to capture common action languages. We show this by providing a translation from action descriptions in action language \mathcal{BC} (Lee et al. 2013) into present-centered temporal logic programs in TEL_f . To this end, we need the following definitions describing the major concepts of \mathcal{BC} . An *action signature* in \mathcal{BC} includes two finite set of symbols, the set \mathbf{F} of fluent constants and the set \mathbf{A} of action constants. Fluent constants are further divided into regular and statically determined fluent constants. A finite set with at least two elements, called *domain*, is assigned to each fluent constant. The set of all values of all domains is denoted by \mathbf{V} . In \mathcal{BC} , an atom is an expression of the form $f = v$, where f is a fluent constant, and v is an element of its domain. If the domain of f consists of truth values \top and \perp , we say that f is Boolean. A *static law* is an expression of the form

$$a_0 \text{ if } a_1, \dots, a_m \text{ ifcons } a_{m+1}, \dots, a_n \quad \text{with } n \geq m \geq 0, \tag{55}$$

where each a_i is an atom for $0 \leq i \leq n$. A *dynamic law* is an expression of the form

$$a_0 \text{ after } a_1, \dots, a_m \text{ ifcons } a_{m+1}, \dots, a_n \quad \text{with } n \geq m \geq 0, \tag{56}$$

where a_0 is an atom containing a regular fluent constant, each of a_1, \dots, a_m is an atom or an action constant, and a_{m+1}, \dots, a_n are atoms. Note that statically determined fluent constants are not allowed in the heads of dynamic laws, but only in the heads of static laws (hence the term ‘‘statically determined’’). An *action description* in \mathcal{BC} consists of a finite set of static and dynamic laws. The semantics of action descriptions in \mathcal{BC} is given by a translation into a sequence of logic programs with nested expressions.

More precisely, given an action description D , a sequence of logic programs with nested expressions $(N_i(D))_{i \geq 0}$ is defined as follows. For $l \geq 0$, the signature of each program $N_l(D)$ is given by

$$\{i : (v = f) \mid f \in \mathbf{F}, v \in \mathbf{V}, 0 \leq i \leq l\} \cup \{i : a \mid a \in \mathbf{A}, 0 \leq i \leq l\}.$$

The logic program $N_l(D)$ contains

- for each static law from D of form (55) and for all $0 \leq i \leq l$, the rule

$$i : a_0 \leftarrow i : a_1 \wedge \dots \wedge i : a_m \wedge \neg(i : a_{m+1}) \wedge \dots \wedge \neg(i : a_n),$$

- for each dynamic law from D of the form (56) and for all $0 \leq i < l$, the rule

$$(i + 1) : a_0 \leftarrow i : a_1 \wedge \dots \wedge i : a_m \wedge \neg((i + 1) : a_{m+1}) \wedge \dots \wedge \neg((i + 1) : a_n),$$

- the choice rule $\{0 : a\}$ for every atom a containing a regular fluent constant,

- the choice rule $\{i : a\}$ for every action constant a and every $i < l$,
- for every fluent constant f and every $i \leq l$, where v_1, \dots, v_k are all elements of the domain of f , the rule

$$\leftarrow \neg i : (f = v_1) \wedge \dots \wedge \neg i : (f = v_k),$$

ensuring the existence of values,

- for every fluent constant f , every pair of distinct elements v, w of its domain, and every $i \leq l$, the rule

$$\leftarrow i : (f = v) \wedge i : (f = w),$$

ensuring the uniqueness of values.

The transition system $T(D)$ represented by the action description D is then defined as follows. For every stable model X of $N_0(D)$, the set of atoms a such that $0 : a$ belongs to X is a state of $T(D)$. For every state s and every fluent constant f there is exactly one v such that $f = v$ belongs to the state s . This v is the value of f in state s . For every stable model X of $N_1(D)$, $T(D)$ includes the transition $\langle s_0, \alpha, s_1 \rangle$, where s_i is the set of atoms a such that $i : a$ belongs to X for $i \in \{0, 1\}$, and α is the set of action constants a such that $0 : a$ belongs to X .

We now provide a translation from the nested logic programs defining the semantics of \mathcal{BC} into present-centered temporal logic programs. For any action description D in \mathcal{BC} , we define the temporal logic program $P(D)$ consisting of

- for each static law from D of form (55), the rules

$$a_1 \wedge \dots \wedge a_m \rightarrow a_0 \vee \neg a_{m+1} \vee \dots \vee \neg a_n \quad \text{and} \\ \widehat{\square}(a_1 \wedge \dots \wedge a_m \rightarrow a_0 \vee \neg a_{m+1} \vee \dots \vee \neg a_n),$$

- for each dynamic law from D of form (56), the rule

$$\widehat{\square}(\bullet a_1 \wedge \dots \wedge \bullet a_m \rightarrow a_0 \vee \neg a_{m+1} \vee \dots \vee \neg a_n),$$

- the formula $a \vee \neg a$ for every atom a containing a regular fluent constant,
- the formulas $a \vee \neg a$ and $\widehat{\square}(a \vee \neg a)$ for every action constant a ,
- for every fluent constant f , where v_1, \dots, v_k are all elements of the domain of f , the rules

$$(f = v_1) \vee \dots \vee (f = v_k) \quad \text{and} \\ \widehat{\square}((f = v_1) \vee \dots \vee (f = v_k)),$$

- for every fluent constant f and every pair of distinct elements v, w of its domain, the rules

$$(f = v) \wedge (f = w) \rightarrow \perp \quad \text{and} \\ \widehat{\square}((f = v) \wedge (f = w) \rightarrow \perp).$$

With this, we get the following result.

Theorem 17 (Relation between stable models of $N_l(D)$ and $P(D)$)

Let D be an action description in \mathcal{BC} over action signature $\langle \mathbf{A}, \mathbf{F}, \mathbf{V} \rangle$, consisting of actions \mathbf{A} , fluents \mathbf{F} and values \mathbf{V} .

Then, we have the following:

1. If $(X_i)_{i \in [0..l]}$ is a temporal stable model of $P(D)$ of length $l > 0$, then

$$\{i : a \mid a \in X_i, a \in \mathbf{A}, 0 \leq i < l\} \cup \{i : (f = v) \mid (f = v) \in X_i, f \in \mathbf{F}, v \in \mathbf{V}, 0 \leq i < l\}$$

is a stable model of $N_{l-1}(D)$.

2. If X is a stable model of $N_l(D)$ with $l \geq 0$, then

$$(\{(v = f) \mid (i : (v = f)) \in X, v \in \mathbf{V}, f \in \mathbf{F}\} \cup \{a \mid (i : a) \in X, a \in \mathbf{A}\})_{0 \leq i \leq l},$$

is a temporal stable model of $P(D)$.

This provides us with an alternative characterization of the transition system corresponding to an action description in \mathcal{BC} .

Corollary 7

Let D be an action description in \mathcal{BC} and $P(D)$ the logic program defined as above. Then,

1. for every stable model $(X_i)_{i \in [0..1]}$ of $P(D)$ of length one, the set of atoms that belong to X_0 is a state of the transition system $T(D)$ and
2. for every stable model $(X_i)_{i \in [0..2]}$ of $P(D)$ of length two, the transition system $T(D)$ includes the transition $\langle s_0, \alpha, s_1 \rangle$, where

- (a) s_i is the set of atoms a such that a belongs to X_i for $i \in \{0, 1\}$ and
- (b) α is the set of action constants that belong to X_0 .

Lee *et al.* (2013) showed that paths of length l in the transition system described by an action description D correspond to the stable models of $N_l(D)$. Using Theorem 17, we can also characterize these paths in terms of temporal stable models.

Corollary 8 (Translation from \mathcal{BC} into TEL_f)

Let D be an action description in \mathcal{BC} and $P(D)$ the logic program defined as above.

A temporal trace $(X_i)_{i \in [0..l]}$ is a temporal stable model of $P(D)$ iff the sequence $(s_i)_{i \in [0..l]}$ where s_i is the set of atoms in X_i is a path in the transition system corresponding to D .

7 Related work

Covering the vast literature that falls in the intersection of temporal reasoning and logic programming is too ambitious for the purpose of the current survey (focused on extending ASP with modal, LTL operators) and would surely require a more extensive document. In this section, we overview those approaches that show a closer connection to the kind of temporal logic programs considered in the paper. We divide this related work into two categories: (1) extending the language with modal operators based on linear traces and (2) using logic programming for temporal reasoning.

7.1 Extending the syntax with linear modalities

As mentioned in the introduction, the first use of the term “Temporal Logic Programming” (Abadi and Manna 1989) dates back to the late 1980s, appearing soon after

the introduction of modal extensions of Prolog (Fariñas del Cerro 1986; Bieber et al. 1988). Several Logic Programming (LP) languages dealing with LTL operators were proposed (Moszkowski 1986; Fujita et al. 1986; Gabbay 1987b; Abadi and Manna 1989; Orgun and Wadge 1990; Baudinet 1992). The latter, a formalism called TEMPLOG, is perhaps a prominent case from a logical point of view. It provides a logical semantics in terms of a least LTL-model, in the spirit of the well-known least Herbrand model for positive¹⁴ logic programs (van Emden and Kowalski 1976). (Cabalar et al. 2015, Theorem 5) proved that TEMPLOG is actually subsumed by TEL, that is, the latter can be used as a generalization of the former for an arbitrary temporal syntax that includes default negation. However, although TEL provides a common underlying semantics to TEMPLOG and our temporal programs (which are the basis of `telingo`), these two languages understand the (modal) temporal logic programming paradigm in a substantially different way, analogous to the differences between Prolog and ASP. In particular, TEMPLOG understands rules in a top-down fashion where the head is considered a *goal* and the body is seen as a *method* (list of subgoals) to achieve the goal. As a result, (future-time) LTL operators are used to represent temporal relations affecting the achievement of these subgoals. As an example, one possible TEMPLOG rule could look like

$$\Box(\textit{printorder} \leftarrow \textit{ack}, \circ(\textit{print}, \Diamond \textit{finished})),$$

meaning that, at any moment, a *printorder* is fulfilled by immediately sending an acknowledgment *ack*, then starting the printer *print* and eventually sending a *finished* message. This TEMPLOG rule naturally corresponds to the TEL implication

$$\Box(\textit{ack} \wedge \circ(\textit{print} \wedge \Diamond \textit{finished}) \rightarrow \textit{printorder}),$$

and has indeed the same semantics in TEL. However, it does not fit into the past–future syntactic fragment used in `telingo`, since the *printorder* in the head is in the relative past of the conditions required in the body, talking about future satisfaction of *print* or *finished*. Following the ASP bottom-up understanding of rules, which is closer to causal laws in action languages, this same example would be represented instead as the past–future formula:

$$\Box(\textit{printorder} \rightarrow \textit{ack} \wedge \circ(\textit{print} \wedge \Diamond \textit{finished})),$$

which can be reduced afterwards to a present-centered logic program. To sum up, TEL is rich enough to cover both bottom-up and top-down readings of program rules, but for its use for temporal ASP, it suffices with a syntactic fragment where past is checked in the bodies and future is used in the head. One last LTL-based LP extension was the preliminary definition of *temporal answer sets* provided by Cabalar (1999). The main difference of that approach with respect to TEL is that the former relied on a different three-valued semantics that was not a proper extension of Equilibrium Logic, losing some of the interesting properties that the latter has shown as a logical foundation of ASP.

Apart from these LP extensions strictly based on LTL, there exist other approaches that introduced other temporal modalities in LP or in ASP and relied on linear-time interpretations or traces. The most relevant and closely related approach is the one by Giordano et al. (2013), where a different definition of *temporal answer set* is provided.

¹⁴ Note that the different semantics for negation as failure were still in their early steps at that moment.

In this case, rather than LTL, the temporal approach used as a basis was *Dynamic Linear Temporal Logic* or DLTL (Henriksen and Thiagarajan 1999), a linear-time variant of the well-known *Dynamic Logic* (Harel et al. 2000). Besides, that definition of temporal answer sets was only applicable to a syntactic fragment of DLTL (also called there “temporal logic programs”) since the semantics relied on a temporal extension of the classical (syntactic) *reduct* transformation (Gelfond and Lifschitz 1988). In order to facilitate a formal comparison, Aguado et al. (2013) introduced a proper extension of TEL that covers the full syntax of DLTL. The main result in that paper proved that, in fact, this extension of TEL for DLTL subsumes the semantics by Giordano et al. (2013) and, in fact, provides its generalization for DLTL arbitrary formulas, without depending on syntactic transformations.

Thanks to the incorporation of regular expressions, dynamic modalities are very interesting for the specification of *control rules* describing steps that are required to be followed by any solution to a temporal problem. Differently from DLTL, another (perhaps more direct) approach to introducing dynamic logic modalities on top of a linear-time semantics is the so-called *Linear Dynamic Logic* or LDL (De Giacomo and Vardi 2013). Both DLTL and LDL are more expressive than LTL, while satisfiability in the three logics share the same PSPACE complexity. However, DLTL uses regular expressions as modifiers of the *until* and *release* operators, whereas the syntax of LDL is closer to the usual in dynamic logic, where modalities include the standard necessity $[\rho]\varphi$ and possibility $\langle\rho\rangle\varphi$ constructs, where ρ is a regular expression. Besides, DLTL regular expressions are built on a separate signature (a set of atomic *actions*) and do not accept the test construct $\varphi?$ from dynamic logic. LTL operators can be encoded both in LDL and DLTL, although the encodings for the latter require an explicit use of the (finite) set of atomic actions. One more advantage of LDL is that a finite trace variant LDL_f has also been proposed and that both LDL and LDL_f properly generalize LTL and LTL_f respectively. All these reasons make LDL a more promising candidate to incorporate dynamic operators in ASP. As a result, Bosser et al. (2018) followed analogous steps to TEL with respect to LTL and introduced a non-monotonic extension of LDL called *Linear Dynamic Equilibrium Logic* (DEL). Moreover, a first implementation of LDL operators in `telingo` has been recently presented by Cabalar et al. (2020), even though they are only allowed in integrity constraints. Although this restriction will be lifted in the future, it already supports an agreeable modeling methodology for dynamic domains separating action and control theories. The idea is to model the actual action theory with temporal rules, fixing static and dynamic laws, while the control theory, enforcing certain (sub)trajectories, is expressed by integrity constraints using dynamic formulas. This is similar to the pairing of action theories in situation calculus and Golog programs (Levesque et al. 1997).

Other modal temporal extensions of LP that rely on linear time have been recently introduced for *stream* reasoning, that is, reactive reasoning over time-varying data sequences. An approach closely related to TEL is LARS (Beck et al. 2015) that deals with time windows or intervals and combines modal LTL operators \square and \diamond with metric and time-point based connectives. Although LARS’ semantics is also based on discrete linear-time, an important difference is that its traces are organized in intervals of time points called *streams*. An informal discussion comparing TEL and LARS was done by Beck et al. (2016) although a formal connection to the corresponding metric extension of TEL (Cabalar et al. 2020) remains to be explored. As proved by Kamp (1968), the

use of discrete time in LTL is not a limitation because any Dedekind complete ordering structure can be used as a timeline, and this includes natural numbers \mathbb{N} or integer numbers \mathbb{Z} but also real \mathbb{R} numbers¹⁵. However, when metric operators come into play, the choice of *dense* time actually leads to a richer expressiveness, usually at the price of undecidability for an arbitrary syntax. For this reason, the most common approach for dense metric extensions is to identify syntactic fragments that keep amenable computational properties, like the case of the so-called *bounded universal Horn formulas* (Brzoska 1995) for positive programs. A more general fragment extending LP with metric modalities is DatalogMTL, a temporal metric extension of Datalog presented by Wałęga et al. (2019) and with applications to stream reasoning. Recently, this approach has been extended to programs with negation (Wałęga et al. 2020) but forbidding the use of \diamond , \mathbb{U} or \mathbf{S} in rule heads. This fragment is decidable and keeps favorable computational properties for the case of an integer timeline, although it jumps to undecidability for dense time, even if considering the rational numbers. Again, a formal connection to the metric extension of TEL (Cabalar et al. 2020) remains to be studied.

The combination of temporal modalities and non-monotonic reasoning is not exclusive to logic programming approaches and was also explored inside the area of Reasoning about Actions and Change. An early approach using LTL is, for instance, the *Past Temporal Logic for Actions* or PTLA (Mendez et al. 1996), an action language incorporating past LTL operators in the conditions of action laws. These action descriptions were then translated into logic programs under the ASP semantics using a similar transformation as the one presented in Section 5.1. Several authors have explored the formalization of action and change using Dynamic Logic (Baldoni et al. 1997; Schwind 1997; Castilho et al. 1999) or DLTL (Giordano et al. 2001), though only the first one relies on a logic programming paradigm.

Another prominent AI field for reasoning about dynamic systems where temporal logic has played an important role is Planning. In this area, the system behavior is specified in terms of some planning-specific language like STRIPS or PDDL, with a carefully limited syntax that avoids the need for an explicit representation of the (non-monotonic) law of inertia. This limited syntax restricts the specification of dynamic systems to a less expressive language than temporal ASP with *telingo*, where defaults, induction or aggregates can be freely combined with temporal operators, but has the advantage of allowing the design and implementation of efficient planning algorithms. Rather than in the description of the transition system, the introduction of temporal formulas in planning has been traditionally related to a richer specification of the *goal*, so that it not only provides a condition about the final state to be reached, but is also extended with temporal formulas (Bacchus and Kabanza 2000; Pistore and Traverso 2001; Bertoli et al. 2001; Kvarnström et al. 2008) imposing constraints on the sequence of actions that form the plan, something that can be used to improve the efficiency of the planning algorithm. This strategy can also be extrapolated to temporal ASP, although the effect of temporal constraints in *telingo* has a different impact on efficiency, given that the algorithm is based on incremental solving. A different strategy for temporal ASP closer to classical planning algorithms was explored in the prototype presented by Cabalar et al. (2019). In

¹⁵ Other dense structures that are not Dedekind complete, like the rational numbers \mathbb{Q} , can be covered by extending the language with Stavi connectives (Gabbay et al. 1980), while keeping decidability.

this case, a classical graph search algorithm was implemented and multi-shot ASP solving was used to compute successor states during search. This has the advantage of being able to explore the whole state space and deciding whether a given planning problem has no solution, something impossible by the incremental horizon strategy followed in `telingo`. However, this planning based prototype had a rather limited syntax and did not allow the use of temporal expressions to incorporate temporal goals.

Other AI areas where temporal logic has been successfully applied are, for instance, Multi Agent Systems (MAS) and Ontologies. In the case of MAS, a prominent example is the system Concurrent `MetateM` (Fisher 1994), a language using LTL operators that exploits Gabbay's separation theorem (Gabbay 1987a) asserting that LTL descriptions can be reduced to a set of implications with the form "*past* \rightarrow *future*." Note that we also exploited this past-future form in `telingo` and our Theorem 16. The main difference between `MetateM` and temporal ASP, however, is that the former relies on monotonic LTL, so it does not allow for closed world reasoning, defaults or induction. On the other hand, `MetateM` incorporates high level constructs for agents specification and message passing that are not present in `telingo`.

Finally, in the case of LTL in Ontologies, the combination of Description Logics (Baader *et al.* 2003) with temporal patterns is an important field of knowledge representation that has been widely studied in the literature (see, for instance, the surveys by Artale and Franconi (2000), Artale and Franconi (2005) and Lutz *et al.* (2008)). However, as happened in the MAS case, the result of these combinations are also monotonic. Combining ontological reasoning and logic programming has also been extensively studied, with the ASP extension called *Hybrid Knowledge Bases* (de Bruijn *et al.* 2010) being perhaps the closest one to our approach, as it is based on Quantified Equilibrium Logic. In a preliminary work, Cabalar and Schaub (2019) presented a combination of hybrid knowledge bases with the temporal ontology framework ALC-LTL by Baader *et al.* (2008). This allowed the use of temporal ontological axioms in temporal logic programs, but no implementation has been made so far.

7.2 Using logic programming for temporal reasoning

A second line of related approaches are those that perform temporal reasoning *using* logic programming (especially Datalog or ASP) as a reasoning engine. In this case, the idea is not necessarily to extend the logic programming syntax with temporal modalities, but to use logic programming as a tool for query answering or model checking with respect to some standard, monotonic temporal logic (most frequently, LTL). If we use Datalog or ASP, the main inconvenience appears when trying to reason about traces of infinite length, since these formalisms historically required a finite domain to allow for a previous grounding of the logic program. For instance, the use of functions was usually forbidden (or limited) to keep a finite Herbrand universe. This limitation was overcome by the early proposal by Chomicki and Imielinski (1988) consisting in the extension of Datalog with a unary *successor* function, to be used inside a fixed parameter in predicates. This approach, called `Datalog1S`, allows for reasoning about infinite traces while keeping decidability. Moreover, it was proved (Baudinet *et al.* 1993) to be equally expressive to `TEMPLOG`, although both languages treat time in a different way: `TEMPLOG` uses *intensional*, temporal modalities, whereas `Datalog1S` uses an *extensional* approach,

representing time points as explicit parameters. Later on, this formalism was extended to so-called Datalog_{NS} (Chomicki and Imielinski 1993) and used as a reasoning system for answering temporal queries expressed in past LTL. A similar approach to Datalog_{IS}, based on the introduction of a successor unary function, was adopted by Eiter and Šimkus (2009) for the more general case of ASP. Other extensions of Datalog that have been used as reasoning engines for temporal logics are Datalog LITE (Gottlob et al. 2002) or *inf*-Datalog (Guessarian et al. 2003). Notice that all these cases are variants of Datalog used to perform model checking for temporal logics such as LTL, CTL or μ -calculus, but are not modal temporal extensions of Datalog.

Another extensional alternative is, rather than using a successor function, to treat time points directly as integer numbers, relying on Constraint Logic Programming (CLP) to solve the integer constraints imposed by temporal expressions. This was the approach adopted by Brzoska (1991), who emulated the operational behavior of TEMPLOG using a CLP translation where each predicate received an additional integer argument, representing time. One more extensional temporal approach based on CLP is, for instance, the Temporal Annotated CLP formalism (Frühwirth 1996). In this case, predicates can be annotated with expressions such as “at T ”, “th I ” or “in I ”, respectively meaning “at time point T ”, “throughout interval I ” and “in interval I ”, respectively. Finally, an early approach using CLP for translating temporal formulas was introduced by Hrycej (1988), but in this case, rather than treating time points as integers, the constraint solver used qualitative constraints for Allen’s Interval Calculus (Allen 1983).

A prominent application of ASP for temporal reasoning is the use of ASP solvers to perform bounded model checking of temporal formulas. The technique of *bounded model checking* (BMC) was introduced by Biere et al. (1999) as a variant of model checking that looks for countermodels of a temporal formula, but only for traces up to some bounded length $k \in \mathbb{N}$. The original approach translated the LTL formula into a propositional theory and then used a SAT solver as a backend. Heljanko and Niemelä (2003) proved that an ASP solver could be used instead of SAT, providing a more compact translation of LTL formulas while keeping an efficient implementation. A similar approach was later followed by Giordano et al. 2012; 2013 to implement BMC for action theories and temporal answer sets. Again, although the use of ASP for BMC is clearly in the intersection of temporal logics and logic programming, it does not necessarily represent a temporal extension of the latter. Note, for instance, that the tool by Heljanko and Niemelä (2003) provides a model checker whose input language are temporal formulas with the standard *monotonic* semantics of LTL.

Finally, one more application of ASP for temporal reasoning is its use for encoding Qualitative (Spatio-)Temporal Calculi (Li 2012; Brenton et al. 2016). In this case, the target temporal formalism is not a modal temporal logic, but some calculus for a set of qualitative relations like, for instance, the aforementioned Allen’s Interval Calculus. For this reason, the relation to the approach studied in the current survey is weaker while, once again, ASP is used as a reasoning engine for encoding the temporal calculi.

8 Conclusion

We have provided a wide overview of the main definitions and recent results for the formalism of TEL, a combination of Equilibrium Logic (a logical characterization of ASP)

with Linear-Time Temporal Logic. After more than a decade of study, the knowledge about TEL has achieved a high degree of maturity, both at a fundamental level and also at an incipient practical application. An important breakthrough for the latter has been the introduction of a finite trace variant, TEL_f , more aligned with the usual problem solving orientation followed in ASP. This has opened the possibility of introducing temporal operators in ASP solving with the construction of the tool `telingo`, a temporal extension of `clingo` that exploits the incremental solving capabilities of the latter. Moreover, the definition of TEL_f has also provided a methodology for the study of other extensions, like the introduction of dynamic logic operators (Cabalar *et al.* 2020) or, more recently, the definition of temporal metric constructions (Cabalar *et al.* 2020) that are planned to be further developed and combined with other features inside the software packages of potassco.org.

This survey has been focused on the LTL extension of ASP where many interesting lines remain to be explored yet. For instance, we still miss an axiomatization that covers the finite trace variant of the monotonic basis, THT_f , or a general axiomatization that is applicable to THT , that is, regardless of the trace length. Also, properties related to inter-definability of operators need to be further explored for the finite trace case. Related to this, Aguado *et al.* (2020) proposed the introduction of an explicit negation operator that allows for inter-definability of temporal operators like *until* and *release* using De Morgan laws. The connection of TEL_f to finite automata is also a topic of active current research. Automata construction techniques may become crucial in the future both for formal verification of dynamic systems specified with `telingo` but also for exploiting temporal expressions in a more compact way during computation, rather than making a full unfolding into regular ASP programs. Another important line of future study is how to exploit temporal constructions for grounding: right now, `telingo` does not recognize temporal information at that stage. A first study about temporal grounding was provided by Aguado *et al.* (2017), but it focused on infinite traces, making these results not extrapolatable to the case of `telingo`, where the hypothesis of finite traces can be exploited more efficiently.

In a broader picture, we are facing a true rebirth of (Modal) Temporal Logic Programming, which is becoming a thrilling area with an intense research both on foundations and in practical implementation. As we have seen in Section 7, there exist several temporal extensions of ASP and Datalog with applications to planning and problem solving on dynamic domains, temporal Knowledge Representation and, more recently, stream reasoning. It seems clear that practical applications may require a richer language than the LTL syntax, as happens with some of the ASP extensions introducing metric constructs or path expressions from dynamic logic. One advantage of the approach we have presented is its characterization in terms of Equilibrium Logic, allowing for an easy adaptation to other modal temporal logics (first extend the modal logic with HT and then add the equilibrium models condition). However, there exist many alternative temporal approaches and choosing which one is the most suited (in terms of expressiveness, complexity and ease of use) for a dynamic problem at hand may turn out to be a challenging question in the future.

Acknowledgments

We are thankful to the anonymous reviewers for their thorough work and their useful suggestions that have helped to improve the paper. A special thanks goes to Mirosław Trzuszczński for his support in improving the quality of our paper. We are especially grateful to David Pearce, whose help and collaboration on Equilibrium Logic was the seed for a great part of the current paper. This work was partially supported by MICINN, Spain, grant PID2020-116201GB-I00, Xunta de Galicia, Spain (GPC ED431B 2019/03), Région Pays de la Loire, France, (projects EL4HC and étoiles montantes CTASP), European Union COST action CA-17124, and DFG grants SCHA 550/11 and 15, Germany.

References

- ABADI, M. AND MANNA, Z. 1989. Temporal logic programming. *Journal of Symbolic Computation* 8, 277–295.
- AGUADO, F., CABALAR, P., DIÉGUEZ, M., PÉREZ, G. AND VIDAL, C. 2013. Temporal equilibrium logic: A survey. *Journal of Applied Non-Classical Logics* 23, 1–2, 2–24.
- AGUADO, F., CABALAR, P., DIÉGUEZ, M., PÉREZ, G. AND VIDAL, C. 2017. Temporal equilibrium logic with past operators. *Journal of Applied Non-Classical Logics* 27, 3–4, 161–177.
- AGUADO, F., CABALAR, P., FANDINNO, J., PÉREZ, G. AND VIDAL, C. 2020. Explicit negation in linear-dynamic equilibrium logic. In *Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20)*, G. De Giacomo, A. Catalá, B. Dilkina, M. Milano, S. Barro, A. Bugarín and J. Lang, Eds. IOS Press, 569–576.
- AGUADO, F., CABALAR, P., PEARCE, D., PÉREZ, G. AND VIDAL, C. 2015. A denotational semantics for equilibrium logic. *Theory and Practice of Logic Programming* 15, 4–5, 620–634.
- AGUADO, F., CABALAR, P., PÉREZ, G., VIDAL, C. AND DIÉGUEZ, M. 2017. Temporal logic programs with variables. *Theory and Practice of Logic Programming* 17, 2, 226–243.
- AGUADO, F., PÉREZ, G. AND VIDAL, C. 2013. Integrating temporal extensions of answer set programming. In *Proceedings of the Twelfth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*, P. Cabalar and T. Son, Eds. Lecture Notes in Artificial Intelligence, vol. 8148. Springer-Verlag, 23–35.
- ALLEN, J. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM* 26, 11, 832–843.
- ARTALE, A. AND FRANCONI, E. 2000. A survey of temporal extensions of description logics. *Annals of Mathematics and Artificial Intelligence* 30, 1–4, 171–210.
- ARTALE, A. AND FRANCONI, E. 2005. Temporal description logics. In *Handbook of Temporal Reasoning in Artificial Intelligence*, M. Fisher, D. Gabbay and L. Vila, Eds. Foundations of Artificial Intelligence, vol. 1. Elsevier Science, 375–388.
- BAADER, F., CALVANESE, D., MCGUINNESS, D., NARDI, D. AND PATEL-SCHNEIDER, P., Eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.
- BAADER, F., GHILARDI, S. AND LUTZ, C. 2008. LTL over description logic axioms. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, G. Brewka and J. Lang, Eds. AAAI Press, 684–694.
- BACCHUS, F. AND KABANZA, F. 2000. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence* 116, 1–2, 123–191.
- BALBIANI, P., BOUDOU, J., DIÉGUEZ, M. AND FERNÁNDEZ-DUQUE, D. 2020. Intuitionistic linear temporal logics. *ACM Transactions on Computational Logic* 21, 2, 14:1–14:32.
- BALBIANI, P. AND DIÉGUEZ, M. 2016. Temporal here and there. In *Proceedings of the Fifteenth*

- European Conference on Logics in Artificial Intelligence (JELIA'16)*, L. Michael and A. Kakas, Eds. Lecture Notes in Artificial Intelligence, vol. 10021. Springer-Verlag, 81–96.
- BALDONI, M., GIORDANO, L., MARTELLI, A. AND PATTI, V. 1997. An abductive proof procedure for reasoning about actions in modal logic programming. In *Proceedings of the Sixth International Workshop on Non-Monotonic Extensions of Logic Programming (NMELP'96)*, J. Dix, L. Pereira, and T. Przymusiński, Eds. Lecture Notes in Computer Science, vol. 1216. Springer-Verlag, 132–150.
- BAUDINET, M. 1992. A simple proof of the completeness of temporal logic programming. In L. Fariñas del Cerro and M. Penttonen, Eds. Clarendon Press, 51–83.
- BAUDINET, M., CHOMICKI, J. AND WOLPER, P. 1993. Temporal deductive databases. In *Temporal Databases: Theory, Design, and Implementation*, A. Tansel, J. Clifford, S. Gadia, A. Segev and R. Snodgrass, Eds. Benjamin/Cummings, 294–320.
- BECK, H., DAO-TRAN, M. AND EITER, T. 2016. Equivalent stream reasoning programs. In *Proceedings of the Twenty-fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*, R. Kambhampati, Ed. IJCAI/AAAI Press, 929–935.
- BECK, H., DAO-TRAN, M., EITER, T. AND FINK, M. 2015. LARS: A logic-based framework for analyzing reasoning over streams. In *Proceedings of the Twenty-Ninth National Conference on Artificial Intelligence (AAAI'15)*, B. Bonet and S. Koenig, Eds. AAAI Press, 1431–1438.
- BERTOLI, P., CIMATTI, A., ROVERI, M. AND TRAVERSO, P. 2001. Planning in nondeterministic domains under partial observability via symbolic model checking. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01)*, B. Nebel, Ed. Morgan Kaufmann Publishers, 473–478.
- BIEBER, P., FARIÑAS DEL CERRO, L. AND HERZIG, A. 1988. MOLOG: a modal PROLOG. In *Proceedings of the Ninth International Conference on Automated Deduction (CADE'88)*, E. Lusk and R. Overbeek, Eds. Lecture Notes in Computer Science, vol. 310. Springer-Verlag, 762–763.
- BIERE, A., CIMATTI, A., CLARKE, E. AND ZHU, Y. 1999. Symbolic model checking without BDDs. In *Proceedings of the Fifth International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS'99)*, R. Cleaveland, Ed. Lecture Notes in Computer Science, vol. 1579. Springer-Verlag, 193–207.
- BOSSER, A., CABALAR, P., DIÉGUEZ, M. AND SCHAUB, T. 2018. Introducing temporal stable models for linear dynamic logic. In *Proceedings of the Sixteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'18)*, M. Thielscher, F. Toni, and F. Wolter, Eds. AAAI Press, 12–21.
- BOZZELLI, L., MAUBERT, B. AND PINCHINAT, S. 2015. Unifying hyper and epistemic temporal logics. In *Proceedings of the Eighteenth International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'15)*, A. Pitts, Ed. Lecture Notes in Computer Science, vol. 9034. Springer-Verlag, 167–182.
- BOZZELLI, L. AND PEARCE, D. 2015. On the complexity of temporal equilibrium logic. In *Proceedings of the Thirtieth Annual Symposium on Logic in Computer Science (LICS'15)*. IEEE Computer Society Press, 645–656.
- BRENTON, C., FABER, W. AND BATSAKIS, S. 2016. Answer set programming for qualitative spatio-temporal reasoning: Methods and experiments. In *Technical Communications of the Thirty-second International Conference on Logic Programming (ICLP'16)*, M. Carro and A. King, Eds. OpenAccess Series in Informatics (OASICs), vol. 52. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 4:1–4:15.
- BRZOSKA, C. 1991. Temporal logic programming and its relation to constraint logic programming. In *Proceedings of the International Symposium on Logic Programming*, V. Saraswat and K. Ueda, Eds. MIT Press, 661–677.
- BRZOSKA, C. 1995. Temporal logic programming in dense time. In *Proceedings of the International Symposium on Logic Programming*, J. Lloyd, Ed. MIT Press, 303–317.

- BÜCHI, J. 1962. On a decision method in restricted second order arithmetic. In *Proceedings of the International Congress on Logic, Methodology, and Philosophy of Science*, E. Nagel, P. Suppes, and A. Tarski, Eds. Stanford University Press, 1–11.
- CABALAR, P. 1999. Temporal answer sets. In *Proceedings of the Joint Conference on Declarative Programming (AGP'99)*, M. Meo and M. Ferro, Eds, 351–366.
- CABALAR, P. AND DEMRI, S. 2011. Automata-based computation of temporal equilibrium models. In *Proceedings of the Twenty-first International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR'11)*, G. Vidal, Ed. Lecture Notes in Computer Science, vol. 7225. Springer-Verlag, 57–72.
- CABALAR, P. AND DIÉGUEZ, M. 2011. STELP — A tool for temporal answer set programming. In *Proceedings of the Eleventh International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'11)*, J. Delgrande and W. Faber, Eds. Lecture Notes in Artificial Intelligence, vol. 6645. Springer-Verlag, 370–375.
- CABALAR, P. AND DIÉGUEZ, M. 2014. Strong equivalence of non-monotonic temporal theories. In *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*, C. Baral, G. De Giacomo and T. Eiter, Eds. AAAI Press.
- CABALAR, P., DIÉGUEZ, M., LAFERRIERE, F. AND SCHAUB, T. 2020. Implementing dynamic answer set programming over finite traces. In *Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20)*, G. De Giacomo, A. Catalá, B. Dilkina, M. Milano, S. Barro, A. Bugarín and J. Lang, Eds. IOS Press, 656–663.
- CABALAR, P., DIÉGUEZ, M., SCHAUB, T. AND SCHUHMAN, A. 2020. Towards metric temporal answer set programming. *Theory and Practice of Logic Programming* 20, 5, 783–798.
- CABALAR, P., DIÉGUEZ, M., AND VIDAL, C. 2015. An infinitary encoding of temporal equilibrium logic. *Theory and Practice of Logic Programming* 15, 4–5, 666–680.
- CABALAR, P., KAMINSKI, R., MORKISCH, P. AND SCHAUB, T. 2019. *telingo* = ASP + time. In *Proceedings of the Fifteenth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'19)*, M. Balduccini, Y. Lierler, and S. Woltran, Eds. Lecture Notes in Artificial Intelligence, vol. 11481. Springer-Verlag, 256–269.
- CABALAR, P., KAMINSKI, R., SCHAUB, T. AND SCHUHMAN, A. 2018. Temporal answer set programming on finite traces. *Theory and Practice of Logic Programming* 18, 3–4, 406–420.
- CABALAR, P., REY, M. AND VIDAL, C. 2019. A complete planner for temporal answer set programming. In *Proceedings of the Nineteenth EPIA Conference on Artificial Intelligence, (EPIA'19)*, P. Oliveira, P. Novais, and L. Reis, Eds. Lecture Notes in Computer Science, vol. 11805. Springer, 520–525.
- CABALAR, P. AND SCHAUB, T. 2019. Temporal logic programs with temporal description logic axioms. In *Description Logic, Theory Combination, and All That - Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, C. Lutz, U. Sattler, C. Tinelli, A. Turhan and F. Wolter, Eds. Lecture Notes in Computer Science, vol. 11560. Springer-Verlag, 174–186.
- CABALAR, P. AND VEGA, G. P. 2007. Temporal equilibrium logic: A first approach. In *Proceedings of the Eleventh International Conference on Computer Aided Systems Theory (EUROCAST'17)*, R. Moreno-Díaz, F. Pichler and A. Quesada-Arencibia, Eds. Lecture Notes in Computer Science, vol. 4739. Springer-Verlag, 241–248.
- CAMACHO, A., BAIER, J., MUISE, C. AND MCILRAITH, S. 2018. Finite LTL synthesis as planning. In *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS'18)*, M. de Weerdt, S. Koenig, G. Röger and M. Spaan, Eds. AAAI Press, 29–38.
- CASTILHO, M., GASQUET, O. AND HERZIG, A. 1999. Formalizing action and change in modal logic I: the frame problem. *Journal of Logic and Computation* 9, 5, 701–735.
- CHANDRA, A., KOZEN, D. AND STOCKMEYER, L. 1981. Alternation. *Journal of the ACM* 28, 1, 114–133.

- CHOMICKI, J. AND IMIELINSKI, T. 1988. Temporal deductive databases and infinite objects. In *Proceedings of the Seventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, C. Edmondson-Yurkanan and M. Yannakakis, Eds. ACM Press, 61–73.
- CHOMICKI, J. AND IMIELINSKI, T. 1993. Finite representation of infinite query answers. *ACM Transactions on Database Systems* 18, 2, 181–223.
- DE BRUIJN, J., PEARCE, D., POLLERES, A. AND VALVERDE, A. 2010. A semantical framework for hybrid knowledge bases. *Knowledge Information Systems* 25, 1, 81–104.
- DE GIACOMO, G. AND PERELLI, G. 2021. Behavioral QLTL. *CoRR abs/2102.11184*.
- DE GIACOMO, G. AND VARDI, M. 2013. Linear temporal logic and linear dynamic logic on finite traces. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI'13)*, F. Rossi, Ed. IJCAI/AAAI Press, 854–860.
- DEMRI, S., GORANKO, V. AND LANGE, M. 2016. *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.
- DURET-LUTZ, A., LEWKOWICZ, A., FAUCHILLE, A., MICHAUD, T., RENAULT, E. AND XU, L. 2016. Spot 2.0 – A framework for LTL and ω -automata manipulation. In *Proceedings of Fourteenth International Symposium on Automated Technology for Verification and Analysis (ATVA'16)*, C. Artho, A. Legay and D. Peled, Eds. Lecture Notes in Computer Science, vol. 9938. 122–129.
- EITER, T. AND ŠIMKUS, M. 2009. Bidirectional answer set programs with function symbols. In *Proceedings of the Twenty-first International Joint Conference on Artificial Intelligence (IJCAI'09)*, C. Boutilier, Ed. AAAI/MIT Press, 765–771.
- FAGES, F. 1994. Consistency of Clark's completion and the existence of stable models. *Journal of Methods of Logic in Computer Science* 1, 51–60.
- FARIÑAS DEL CERRO, L. 1986. MOLOG: A system that extends PROLOG with modal logic. *New Generation Computing* 4, 1, 35–50.
- FERRARIS, P., LEE, J. AND LIFSCHITZ, V. 2007. A new perspective on stable models. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, M. Veloso, Ed. AAAI/MIT Press, 372–379.
- FISCHER, M. AND LADNER, R. 1979. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences* 18, 2, 194–211.
- FISHER, M. 1994. A survey of concurrent METATEM – The language and its applications. In *Proceedings of the First International Conference on Temporal Logic (ICTL'94)*, D. Gabbay and H. Ohlbach, Eds. Lecture Notes in Computer Science, vol. 827. Springer-Verlag, 480–505.
- FRENCH, T. AND REYNOLDS, M. 2002. A sound and complete proof system for QPTL. In *Proceedings of the Fourth Conference on Advances in Modal Logic*, P. Balbiani, N. Suzuki, F. Wolter, and M. Zakharyashev, Eds. King's College Publications, 127–148.
- FRÜHWIRTH, T. 1996. Temporal annotated constraint logic programming. *Journal of Symbolic Computation* 22, 5–6, 555–583.
- FUJITA, M., KONO, S., TANAKA, H. AND MOTO-OKA, T. 1986. Tokio: Logic programming language based on temporal logic and its compilation to Prolog. In *Proceedings of the Third International Conference on Logic Programming (ICLP'86)*, E. Shapiro, Ed. Lecture Notes in Computer Science, vol. 225. Springer, 695–709.
- GABBAY, D. 1987a. The declarative past and imperative future: Executable temporal logic for interactive systems. In *Proceedings of the Conference on Temporal Logic in Specification*, B. Banieqbal, H. Barringer and A. Pnueli, Eds. Lecture Notes in Computer Science, vol. 398. Springer-Verlag, 409–448.
- GABBAY, D. 1987b. Modal and temporal logic programming. In *Temporal Logics and their Applications*, A. Galton, Ed. Academic Press, Chapter 6, 197–237.

- GABBAY, D., PNUELI, A., SHELAH, S. AND STAVI, J. 1980. On the temporal basis of fairness. In *Proceedings of the Seventh Symposium on Principles of Programming Languages (POPL'80)*, P. Abrahams, R. Lipton and S. Bourne, Eds. ACM Press, 163–173.
- GEBBER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2019. Multi-shot ASP solving with clingo. *Theory and Practice of Logic Programming* 19, 1, 27–82.
- GELFOND, M. AND LIFSCHITZ, V. 1988. Compiling circumscriptive theories into logic programs. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI'88)*, H. Shrobe, T. Mitchell and R. Smith, Eds. AAAI Press / The MIT Press, 455–459.
- GELFOND, M. AND LIFSCHITZ, V. 1993. Representing action and change by logic programs. *Journal of Logic Programming* 17, 2–4, 301–321.
- GERTH, R., PELED, D., VARDI, M. AND WOLPER, P. 1995. Simple on-the-fly automatic verification of linear temporal logic. In *Proceedings of the Fifteenth IFIP International Symposium on Protocol Specification*, P. Dembinski and M. Sredniawa, Eds. IFIP Conference Proceedings, vol. 38. Chapman & Hall, 3–18.
- GIORDANO, L., MARTELLI, A. AND DUPRÉ, D. T. 2012. Achieving completeness in bounded model checking of action theories in ASP. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'12)*, G. Brewka, T. Eiter and S. McIlraith, Eds. AAAI Press.
- GIORDANO, L., MARTELLI, A. AND THESEIDER DUPRÉ, D. 2013. Reasoning about actions with temporal answer sets. *Theory and Practice of Logic Programming* 13, 2, 201–225.
- GIORDANO, L., MARTELLI, M. AND SCHWIND, C. 2001. Reasoning about actions in dynamic linear time temporal logic. *Logic Journal of the IGPL* 9, 2, 273–288.
- GÖDEL, K. 1932. Zum intuitionistischen Aussagenkalkül. *Anzeiger der Akademie der Wissenschaften in Wien*, 65–66.
- GOLDBLATT, R. 1992. *Logics of Time and Computation*. Number 7 in CSLI Lecture Notes. Center for the Study of Language and Information.
- GOTTLOB, G., GRÄDEL, E. AND VEITH, H. 2002. Datalog LITE: A deductive query language with linear time model checking. *ACM Transactions on Computational Logic* 3, 1, 42–79.
- GUESSARIAN, I., FOUSTOUCOS, E., ANDRONIKOS, T. AND AFRATI, F. 2003. On temporal logic versus datalog. *Theoretical Computer Science* 303, 1, 103–133.
- HAREL, D., TIURYN, J. AND KOZEN, D. 2000. *Dynamic Logic*. MIT Press.
- HELJANKO, K. AND NIEMELÄ, I. 2003. Bounded LTL model checking with stable models. *Theory and Practice of Logic Programming* 3, 4–5, 519–550.
- HENRIKSEN, J., JENSEN, J., JØRGENSEN, M., KLARLUND, N., PAIGE, R., RAUHE, T. AND SANDHOLM, A. 1995. Mona: Monadic second-order logic in practice. In *Proceedings of the First International Workshop on Tools and Algorithms for Construction and Analysis of Systems (TACAS'95)*, E. Brinksma, R. Cleaveland, K. Larsen, T. Margaria and B. Steffen, Eds. Lecture Notes in Computer Science, vol. 1019. Springer-Verlag, 89–110.
- HENRIKSEN, J. AND THIAGARAJAN, P. 1999. Dynamic linear time temporal logic. *Annals Pure and Applied Logic* 96, 1–3, 187–207.
- HEYTING, A. 1930. Die formalen Regeln der intuitionistischen Logik. In *Sitzungsberichte der Preussischen Akademie der Wissenschaften*. Deutsche Akademie der Wissenschaften zu Berlin, 42–56.
- HOSOI, T. 1966. The axiomatization of the intermediate propositional systems s_2 of Gödel. *Journal of the Faculty of Science of the University of Tokyo* 13, 2, 183–187.
- HRYCEJ, T. 1988. Temporal Prolog. In *Proceedings of the Eighth European Conference on Artificial Intelligence (ECAI'92)*, Y. Kodratoff, Ed. Pitman/Morgan Kaufmann, 296–301.
- KAMP, J. 1968. Tense logic and the theory of linear order. Ph.D. thesis, University of California at Los Angeles.

- KVARNSTRÖM, J., HEINTZ, F. AND DOHERTY, P. 2008. A temporal logic-based planning and execution monitoring system. In *Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS'08)*, J. Rintanen, B. Nebel, J. Beck and E. Hansen, Eds. AAAI Press, 198–205.
- LEE, J., LIFSCHITZ, V. AND YANG, F. 2013. Action language BC: Preliminary report. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI'13)*, F. Rossi, Ed. IJCAI/AAAI Press, 983–989.
- LEVESQUE, H., REITER, R., LESPÉRANCE, Y., LIN, F. AND SCHERL, R. 1997. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming* 31, 1–3, 59–83.
- LI, J. 2012. Qualitative spatial and temporal reasoning with answer set programming. In *Proceedings of the Twenty-fourth IEEE International Conference on Tools with Artificial Intelligence (ICTAI'12)*. IEEE Computer Society Press, 603–609.
- LIFSCHITZ, V. 2002. Answer set programming and plan generation. *Artificial Intelligence* 138, 1–2, 39–54.
- LUKASIEWICZ, J. 1941. Die logik und das grundlagenproblem. *Les Entretiens de Zürich sur les Fondaments et la Méthode des Sciences Mathématiques* 12, 6–9, 82–100.
- LUTZ, C., WOLTER, F. AND ZAKHARYASCHEV, M. 2008. Temporal description logics: A survey. In *Proceedings of the fifteenth International Symposium on Temporal Representation and Reasoning (TIME'08)*, S. Demri and C. Jensen, Eds. IEEE Computer Society Press, 3–14.
- MENDEZ, G., LLOPIS, J., LOBO, J. AND BARAL, C. 1996. Temporal logic and reasoning about actions. In *Proceedings of the Third Symposium on Logical Formalizations of Commonsense Reasoning*.
- MINTS, G. 2010. Cut-free formulations for a quantified logic of here and there. *Annals of Pure and Applied Logic* 162, 3, 237–242.
- MOSZKOWSKI, B. 1986. *Executing Temporal Logic Programs*. Cambridge University Press.
- OIKARINEN, E. AND JANHUNEN, T. 2006. Modular equivalence for normal logic programs. In *Proceedings of the Seventeenth European Conference on Artificial Intelligence (ECAI'06)*, G. Brewka, S. Coradeschi, A. Perini and P. Traverso, Eds. IOS Press, 412–416.
- ORGUN, M. AND WADGE, W. 1990. Theory and practice of temporal logic programming. L. Fariñas del Cerro and M. Penttonen, Eds. Clarendon Press, 21–50.
- PEARCE, D. 1997. A new logical characterisation of stable models and answer sets. In *Proceedings of the Sixth International Workshop on Non-Monotonic Extensions of Logic Programming (NMELP'96)*, J. Dix, L. Pereira and T. Przymusiński, Eds. Lecture Notes in Computer Science, vol. 1216. Springer-Verlag, 57–70.
- PEARCE, D., TOMPITS, H. AND WOLTRAN, S. 2001. Encodings for equilibrium logic and logic programs with nested expressions. In *Proceedings of the Tenth Portuguese Conference on Artificial Intelligence (EPIA'01)*, P. Brazdil and A. Jorge, Eds. Lecture Notes in Computer Science, vol. 2258. Springer-Verlag, 306–320.
- PEARCE, D. AND VALVERDE, A. 2008. Quantified equilibrium logic and foundations for answer set programs. In *Proceedings of the Twenty Fourth International Conference of Logic Programming (ICLP'08)*, M. Garcia De La Banda and E. Pontelli, Eds. Lecture Notes in Computer Science, vol. 5366. Springer-Verlag, 546–560.
- PISTORE, M. AND TRAVERSO, P. 2001. Planning as model checking for extended goals in non-deterministic domains. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01)*, B. Nebel, Ed. Morgan Kaufmann Publishers, 479–486.
- PNUELI, A. 1977. The temporal logic of programs. In *Proceedings of the Eight-teenth Symposium on Foundations of Computer Science (FOCS'77)*. IEEE Computer Society Press, 46–57.
- SCHWIND, C. 1997. A logic based framework for action theories. In *Proceedings of the Tbilisi Symposium on Logic, Language and Computation*, J. Ginzburg, Z. Khasidashvili, C. Vogel, J. Lévy and E. Vallduví, Eds. 275–291.

- SIMPSON, A. 1994. The Proof Theory and Semantics of Intuitionistic Modal Logic. Ph.D. thesis, University of Edinburgh.
- SISTLA, A. 1983. Theoretical issues in the design and verification of distributed systems. Ph.D. thesis, Harvard University.
- SISTLA, A. AND CLARKE, E. 1985. The complexity of propositional linear temporal logic. *Journal of the ACM* 32, 3, 733–749.
- SISTLA, A., VARDI, M. AND WOLPER, P. 1987. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science* 49, 217–237.
- VAN DALEN, D. 2001. Intuitionistic logic. In *Handbook of Philosophical Logic*, D. Gabbay and F. Guenther, Eds. Vol. 3. Springer-Verlag, 225–339.
- VAN EMDEN, M. AND KOWALSKI, R. 1976. The semantics of predicate logic as a programming language. *Journal of the ACM* 23, 4, 733–742.
- WALEGA, P., CUENCA GRAU, B., KAMINSKI, M. AND KOSTYLEV, E. 2019. DatalogMTL: Computational complexity and expressive power. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI'19)*, S. Kraus, Ed. ijcai.org, 1886–1892.
- WALEGA, P., TENA CUCALA, D., KOSTYLEV, E. AND CUENCA GRAU, B. 2020. DatalogMTL with negation under stable models semantics. In *Proceedings of the Eighteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'22)*, M. Bienvenu, G. Lakemeyer, and E. Erdem, Eds. AAAI Press, 609–618.
- WOLPER, P. 1983. Temporal logic can be more expressive. *Information and Control* 56, 1/2, 72–99.
- ZHU, S., PU, G. AND VARDI, M. 2019. First-order vs. second-order encodings for LTLf-to-automata translation. In *Proceedings of the Fifteenth Annual Conference on Theory and Applications of Models of Computation (TAMC'19)*, T. Gopal and J. Watada, Eds. Lecture Notes in Computer Science, vol. 11436. Springer-Verlag, 684–705.
- ZHU, S., TABAJARA, L., LI, J., PU, G. AND VARDI, M. 2017. Symbolic LTLf synthesis. In *Proceedings of the Twenty-sixth International Joint Conference on Artificial Intelligence (IJCAI'17)*, C. Sierra, Ed. IJCAI/AAAI Press, 1362–1369.