

LTL on Weighted Finite Traces: Formal Foundations and Algorithms

Carmine Dodaro, Valeria Fionda* and Gianluigi Greco

Department of Mathematics and Computer Science, University of Calabria

{carmine.dodaro, valeria.fionda, gianluigi.greco}@unical.it

Abstract

LTL on finite traces (LTL_f) is a logic that attracted much attention in recent literature, due to its ability to formalize the qualitative behavior of dynamical systems in several application domains. However, its practical usage is still rather limited, as LTL_f cannot deal with any quantitative aspect, such as with the costs of realizing some desired behaviour. The paper fills the gap by proposing a weighting framework for LTL_f encoding such quantitative aspects in the traces over which it is evaluated. The complexity of reasoning problems on *weighted traces* is analyzed and compared to that of standard LTL_f , by considering arbitrary formulas as well as classes of formulas defined in terms of relevant syntactic restrictions. Moreover, a reasoner for LTL_f on weighted traces is presented, and its performance are assessed on benchmark data.

1 Introduction

Linear Temporal Logic over finite traces (short: LTL_f) is a logic formalism based on interpreting standard LTL formulas [Pnueli, 1977] over finite, yet unbounded, successions of states (see, e.g., [De Giacomo and Vardi, 2013]). The formalism has been recently attracting much attention in the AI literature, as it fits several scenarios where the traditional interpretation over infinite traces is not appropriate. For instance, LTL_f has been exploited to formalize and validate the behavior of dynamical systems (e.g., [De Giacomo *et al.*, 2014b; De Giacomo and Vardi, 2015]), to reason about planning domains (e.g. [De Giacomo and Rubin, 2018; Zhu *et al.*, 2020; Xiao *et al.*, 2021]), and to declaratively specify, analyze and monitor business processes (e.g. [van der Aalst *et al.*, 2009; Montali *et al.*, 2010; De Giacomo *et al.*, 2014a]).

As a matter of facts, however, the usage of LTL_f in real-world domains is still rather limited, despite a number of efficient reasoning engines have been already made available (see [Li *et al.*, 2020; Fionda and Greco, 2018]). Indeed, a major limitation of LTL_f is that it cannot go beyond reasoning over qualitative properties of the environment, since it misses

the ability to address certain quantitative aspects that are crucial in practical applications, such as time, probabilities, or costs related to realizing some desired behavior.

Clearly enough, LTL_f inherits this limitation from LTL. In fact, the development of “quantitative” frameworks for LTL has been an active area of research in the last two decades. In particular, LTL extensions have been proposed within settings where constants take values from some given lattice and the semantics of Boolean and temporal operators has been recast as operations defined on that lattice [Kupferman and Lustig, 2007; Droste and Vogler, 2012]. Other works have instead studied LTL over weighted transition systems [Faella *et al.*, 2008; de Alfaro *et al.*, 2004; Droste and Rahonis, 2016; Baier *et al.*, 2014]. And, finally, a number of approaches have been defined to introduce quantitative features suited—for instance—to express a tolerance on the number of time instants where a formula that is always required to hold in the future actually does not hold [Kuperberg, 2014], to specify eventualities with a bounded wait time [Kupferman *et al.*, 2009; Alur *et al.*, 2001], to accommodate averaging [Bouyer *et al.*, 2014] or discounting [Mandralli, 2012; de Alfaro *et al.*, 2005] modalities, and to deal with propositional quality operators [Almagor *et al.*, 2016].

Compared to the above body of literature for LTL, considerably less attention has been paid to study LTL_f in quantitative settings. Indeed, even though we might adapt to LTL_f some of the extensions introduced for LTL (see, e.g., [Camacho *et al.*, 2018b]), the typical usage scenarios of LTL_f call for ad-hoc solutions and approaches, which were missing in the literature so far. In the paper, we fill this gap by proposing and studying a “weighting” framework that is specific for LTL_f and where quantitative aspects are encoded in the finite traces over which formulas are evaluated, rather than on the formulas themselves. Roughly, each variable is associated with a weight, and the weight of a state is given as a function of the weights of the variables evaluating true in it; the goal is, then, to compute an optimal model, i.e., such that the overall weight obtained by suitably aggregating the weights of the various (finitely many) states is minimized. In more details:

- We define the concept of *valuation structures* as mechanisms to flexibly equip each trace with a weight, and we identify and exemplify certain monotonicity properties of that structures guaranteeing that *optimal models* always exist for satisfiable LTL_f formulas (Section 2).

*Contact Author

- We depict a clear picture of the computational costs to be paid for the rich expressiveness of the resulting framework compared to standard LTL on finite traces, by considering arbitrary formulas as well as formulas belonging to some relevant syntactic fragments (Section 3).
- We illustrate the architecture of WELTL, a system prototype implemented to reason about LTL_f over weighted traces exploiting our theoretical findings and a rewriting in terms of *Answer Set Programming (ASP)* [Brewka et al., 2011]; and, finally, we assess and discuss its performance on some benchmark data (Section 4).¹

2 Formal Framework

In this section, we introduce a framework for evaluating LTL_f formulas on weighted finite traces, and we illustrate its salient features by discussing some example applications.

2.1 Basic Notions and Results

LTL on finite traces. Let \mathcal{V} be a universe of propositional variables. An LTL_f formula φ over \mathcal{V} is an expression² built according to the grammar $\varphi ::= \text{true} \mid \text{false} \mid x \mid \neg x \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \mathbf{X}(\varphi) \mid \mathbf{X}_w(\varphi) \mid (\varphi \mathbf{U} \varphi) \mid (\varphi \mathbf{R} \varphi)$, where: $x \in \mathcal{V}$; “ \wedge ”, “ \vee ”, and “ \neg ” are the standard Boolean connectives; and “ \mathbf{X} ” (next), “ \mathbf{X}_w ” (weak next), “ \mathbf{U} ” (until), and “ \mathbf{R} ” (release) are temporal operators. A *finite trace* over \mathcal{V} is a sequence $\pi = \pi_0, \pi_1, \dots, \pi_{n-1}$ such that $\pi_i \subseteq \mathcal{V}$, for each $i \in \{0, \dots, n-1\}$. We say that the formula φ holds in π at the *instant* $i \in \{0, \dots, n-1\}$, denoted by $\pi, i \models \varphi$, by inductively considering *subformulas* as follows:

$$\begin{aligned}
 \pi, i \models y & \quad \text{iff } y \in \pi_i \text{ or } y = \text{true}; \\
 \pi, i \models \neg y & \quad \text{iff } y \notin \pi_i \text{ or } y = \text{false}; \\
 \pi, i \models (\varphi_1 \wedge \varphi_2) & \quad \text{iff } \pi, i \models \varphi_1 \text{ and } \pi, i \models \varphi_2; \\
 \pi, i \models (\varphi_1 \vee \varphi_2) & \quad \text{iff } \pi, i \models \varphi_1 \text{ or } \pi, i \models \varphi_2; \\
 \pi, i \models \mathbf{X}(\varphi') & \quad \text{iff } i < n-1 \text{ and } \pi, i+1 \models \varphi'; \\
 \pi, i \models \mathbf{X}_w(\varphi') & \quad \text{iff } \pi, i \models \mathbf{X}(\varphi') \text{ or } i = n-1; \\
 \pi, i \models (\varphi_1 \mathbf{U} \varphi_2) & \quad \text{iff } \exists j \geq i \text{ such that } \pi, j \models \varphi_2 \text{ and } \forall k \\
 & \quad \text{with } i \leq k < j \text{ we have } \pi, k \models \varphi_1; \\
 \pi, i \models (\varphi_1 \mathbf{R} \varphi_2) & \quad \text{iff } \forall i \leq j < n \text{ it holds } \pi, j \models \varphi_2; \\
 & \quad \text{or } \exists j \geq i \text{ such that } \pi, j \models \varphi_1 \text{ and} \\
 & \quad \forall i \leq k \leq j \text{ we have } \pi, k \models \varphi_2;
 \end{aligned}$$

Whenever $\pi, 0 \models \varphi$ holds, then we say that π is a *model* of φ and we just write $\pi \models \varphi$.

Example 2.1 Consider the LTL_f formula $\varphi = (a \vee \mathbf{X}(b)) \wedge (c \mathbf{U} (d \wedge \mathbf{X}_w(a)))$ defined over the set $\mathcal{V} = \{a, b, c, d\}$ of propositional variables. The finite trace $\pi = \{a, d\}$ (consisting of one time instant only) is a model of φ ; indeed, $\pi, 0 \models a \vee \mathbf{X}(b)$ holds because $a \in \pi_0$, while $\pi, 0 \models c \mathbf{U} (d \wedge \mathbf{X}_w(a))$ holds because $d \in \pi_0$ and there is no subsequent time instant (hence, $\mathbf{X}_w(a)$ holds too). Other models are $\pi' = \{c\}, \{b, d\}, \pi'' = \{c\}, \{c, b\}, \{d\}$, and $\pi''' = \{c\}, \{c, b\}, \{d\}, \{a\}$. Instead, note that the trace $\{b, c\}, \{a\}, \{d\}$ is not a model. \triangleleft

Valuation structures. A *valuation structure* \mathbf{v} on \mathcal{V} is a tuple $\mathbf{v} = \langle \circ, *, w \rangle$, where $w : \mathcal{V} \mapsto \mathbb{Z}$ is a function mapping

¹WELTL, the benchmark data and details on the (theoretical/experimental) results are available as Supplementary Material.

²W.l.o.g., we consider formulas in *negated normal form*, i.e., negations are in front of propositional variables only.

each variable $x \in \mathcal{V}$ to an integer $w(x)$, and where \circ and $*$ are two binary operators defined over \mathbb{Z} , which are moreover commutative and associative. For any set $\{x_1, \dots, x_m\} \subseteq \mathcal{V}$, we define $\mathbf{v}(\{x_1, \dots, x_m\})$ as the value $w(x_1) \circ \dots \circ w(x_m)$. Moreover, for a trace $\pi = \pi_0, \pi_1, \dots, \pi_{n-1}$, its *v-weight* is defined as $\mathbf{v}(\pi) = \mathbf{v}(\pi_0) * \dots * \mathbf{v}(\pi_{n-1})$.

Example 2.2 Consider again the setting of Example 2.1, and the valuation structure $\mathbf{v} = \langle +, \times, w \rangle$ where $w(a) = 6$, $w(b) = 2$, $w(c) = 2$, and $w(d) = 1$. Then, we have $\mathbf{v}(\pi) = \mathbf{v}(\{a, d\}) = w(a) + w(d) = 7$. Furthermore, we have that $\mathbf{v}(\pi') = \mathbf{v}(\pi'_0) \times \mathbf{v}(\pi'_1) = 6$, $\mathbf{v}(\pi'') = \mathbf{v}(\pi''_0) \times \mathbf{v}(\pi''_1) \times \mathbf{v}(\pi''_2) = 8$, and $\mathbf{v}(\pi''') = \mathbf{v}(\pi''') \times w(a) = 48$. \triangleleft

Valuation structures are used to single out the “best” models by just looking for minimizing their associated weights. More formally, any model π^* of φ is called *v-optimal* if there is no model π of φ with $\mathbf{v}(\pi) < \mathbf{v}(\pi^*)$.

Example 2.3 In our running example, we have $\mathbf{v}(\pi') < \mathbf{v}(\pi)$, $\mathbf{v}(\pi') < \mathbf{v}(\pi'')$, and $\mathbf{v}(\pi') < \mathbf{v}(\pi''')$ hold. In fact, π' is an optimal model. In particular, note that the models where c holds at the first time instant, while c plus b hold at the second time instant have v-weight at least $2 \times 4 = 8$. \triangleleft

Monotone valuations. Perhaps not surprisingly, optimal models are not guaranteed to exist, even ifad the LTL_f formula given to hand is *satisfiable* (that is, if admits a model). For instance, consider the structure $\mathbf{v} = \langle +, +, -1 \rangle$ where -1 is the constant function mapping every propositional variable $x \in \mathcal{V}$ to -1 . Then, for the formula $\varphi = a$ and any model π (of whatever length), the trace $\bar{\pi}$ obtained by appending the state $\{a\}$ to π is again a model and $\mathbf{v}(\bar{\pi}) = \mathbf{v}(\pi) - 1 < \mathbf{v}(\pi)$.

Motivated by the above observation, we hereinafter restrict ourselves to valuation structures $\mathbf{v} = \langle \circ, *, w \rangle$, which we call *monotone*, enjoying the following properties:

- (P1) $w(x) \geq 0$, for each $x \in \mathcal{V}$;
- (P2) $w_1 \circ w_2 \geq \max\{w_1, w_2\}$, for each pair $w_1, w_2 \geq 0$; and
- (P3) $w_1 * w_2 \geq \max\{w_1, w_2\}$, for each pair $w_1, w_2 \geq 0$.

The following is a simple consequence of monotonicity.

Lemma 2.4 Assume that $\mathbf{v} = \langle \circ, *, w \rangle$ is a monotone valuation structure, and let π_0, \dots, π_{n-1} be a trace. Then, for each pair $0 \leq i < j \leq n-1$, it holds $\mathbf{v}(\pi_0, \dots, \pi_i) \leq \mathbf{v}(\pi_0, \dots, \pi_j)$.

Theorem 2.5 If $\mathbf{v} = \langle \circ, *, w \rangle$ is a monotone valuation structure and φ is satisfiable, then φ admits a v-optimal model.

Proof (Sketch). Let Π be the set of *irreducible* models of φ , that is, none of their prefixes is a model too. Since φ is satisfiable, $\Pi \neq \emptyset$. Then, let π^* be a model of minimum v-weight over all models in Π . We claim that π^* is v-optimal. Indeed, assume by contradiction that $\pi = \pi_0, \dots, \pi_{n-1}$ is a model with $\mathbf{v}(\pi) < \mathbf{v}(\pi^*)$. By construction of π^* , $\pi \notin \Pi$; hence, there is a time instant $i \in \{0, \dots, n-2\}$ such that the prefix π_0, \dots, π_i belongs to Π , which implies $\mathbf{v}(\pi^*) \leq \mathbf{v}(\pi_0, \dots, \pi_i)$. However, by Lemma 2.4, we know that $\mathbf{v}(\pi) \geq \mathbf{v}(\pi_0, \dots, \pi_i)$, which entails $\mathbf{v}(\pi) \geq \mathbf{v}(\pi^*)$, which is impossible. \square




Figure 1: The planning problem P, three models of φ_P , and their values under the valuations structures built over “+” and “max”.

2.2 Examples of Monotone Valuations

Noticeable monotone valuations can be defined over “+” and “max”, e.g. to encode cost-related planning problems. Consider for instance the planning problem P in Figure 1: a set $\mathcal{F} = \{f_1, \dots, f_8\}$ of fluents is given and, by starting from the scenario where f_1 holds, our goal is to end up with f_8 , by applying, possibly in parallel, some of the available actions in $\mathcal{A} = \{a, b, c, d, e\}$. An example solution is the plan where we first execute a (“consuming” f_1 and “producing” f_2, f_3 , and f_4); then, we execute in parallel actions b, c , and d ; and, finally, we execute e . It is well-known (see, e.g., [De Giacomo and Vardi, 2013]) that planning problems P can be easily encoded via LTL_f formulas φ_P over propositional variables in $\mathcal{A} \cup \mathcal{F}$ so that models for φ_P one-to-one correspond with plans for P—at each time instant, we just keep trace of the fluents that hold and of the actions that are executed (see again Figure 1, reporting the model corresponding to the plan discussed above, plus two further models). We next enrich φ_P with valuations suited to encode cost-related variants.

$\langle +, +, c \rangle$. Assume that each action $x \in \mathcal{A}$ is associated with a cost $c(x)$, and consider the structure $\langle +, +, c \rangle$, where $c(y) = 0$ holds, for each fluent $y \in \mathcal{F}$. Then, $\langle +, +, c \rangle$ -optimal models correspond to plans whose total cost is minimized over all the possible plans. For instance, by assuming that all actions cost 1, every model of φ_P is also an optimal one, since all plans reaching the goal have the same cost (namely, 5)—for they need to execute all given actions.

$\langle +, \max, r \rangle$. Suppose that each action $x \in \mathcal{A}$ is associated with the number $r(x)$ of employees required for its execution. If it is necessary, in time of a pandemic, to limit the number of employees that are simultaneously in the office, then we can consider the valuation structure $\langle +, \max, r \rangle$ (with $r(y) = 0$, for each $y \in \mathcal{F}$). For instance, in the setting of Figure 1, by assuming that all actions require 1 employee only, except d requiring 2 employees, we get that $\langle +, \max, r \rangle$ -optimal models of φ_P correspond to plans where d is executed alone.

$\langle \max, +, t \rangle$. Consider a setting where each action $x \in \mathcal{A}$ is associated with the time $t(x)$ required for its execution. If we are interested in minimizing the total time for plan execution, then we can consider the valuation structure $\langle \max, +, t \rangle$ (with $t(y) = 0$, for each $y \in \mathcal{F}$)—in particular, note that at each step, for the actions that are executed in parallel, we just consider the maximum time over all of them. In our running example, by assuming that every action requires a unit of time, the only $\langle \max, +, t \rangle$ -optimal model of φ_P corresponds to the plan where b, c , and d are executed in parallel.

$\langle \max, \max, t \rangle$. Assume, finally, that we want to fix the duration of the work shift, by computing the minimum amount

of time needed to complete each step. This can be modeled via the structure $\langle \max, \max, t \rangle$, so that the total cost is now the maximum completion time over all the steps. For instance, by assuming again that every action requires a unit of time, every model is an optimal one and its cost is 1.

3 Complexity of LTL_f on Weighted Traces

In this section, we study the computational complexity of reasoning about LTL_f on weighted traces. Formally, we consider the \exists MODEL problem that, given as input an integer $\tau \geq 0$, an LTL_f formula φ , and a structure $\mathbf{v} = \langle \circ, *, w \rangle$, asks whether there is a model π of φ such that $\mathbf{v}(\pi) \leq \tau$.

In the analysis, we assume a standard encoding for the input; in particular, \mathbf{v} is given by explicitly listing all elements in its domain with their associated values, while \circ and $*$ are polynomial-time computable and produce outputs having size polynomial w.r.t. the input size. Hereinafter, the size of φ , i.e., the total number of temporal operators, Boolean connectives and propositional variables it contains, is denoted by $\|\varphi\|$.

3.1 Basic Results: Weighted Transition Graph

To prove our basic complexity results, we introduce the notion of *weighted transition graph*, which we formalize in the spirit of the transition system defined by [Li et al., 2020].

An LTL_f formula is in *next Normal Form* (XNF) if all occurrences of the temporal operators R and U are nested within the scope of some X or X_w operator. Note that there is a linear-time conversion from any formula φ into an equivalent XNF one, denoted as $\text{xnf}(\varphi)$; indeed, we can build $\text{xnf}(\varphi)$ from φ by just applying the following transformations:

- $\varphi_1 \text{ U } \varphi_2$ is replaced by $\varphi_2 \vee (\varphi_1 \wedge \text{X}(\varphi_1 \text{ U } \varphi_2))$;
- $\varphi_1 \text{ R } \varphi_2$ is replaced by $\varphi_2 \wedge (\varphi_1 \vee \text{X}_w(\varphi_1 \text{ R } \varphi_2))$.

Given $\text{xnf}(\varphi)$, we then define $\mathcal{V}_X(\varphi)$ (resp., $\mathcal{V}_{X_w}(\varphi)$) as the set of all its X-subformulas (resp., X_w-subformulas) that are not nested within the scope of some other X or X_w operator; moreover, let $\text{xnf}(\varphi)^P$ denote $\text{xnf}(\varphi)$ viewed as a propositional formula over $\mathcal{V} \cup \mathcal{V}_X(\varphi) \cup \mathcal{V}_{X_w}(\varphi)$.

Example 3.1 For the formula φ of Example 2.1, we have

$$\text{xnf}(\varphi) = (a \vee \text{X}(b)) \wedge ((d \wedge \text{X}_w(a)) \vee (c \wedge \text{X}(c \text{ U } (d \wedge \text{X}_w(a))))).$$

So, $\mathcal{V}_X(\varphi) = \{\text{X}(b), \text{X}(c \text{ U } (d \wedge \text{X}_w(a)))\}$ and $\mathcal{V}_{X_w}(\varphi) = \{\text{X}_w(a)\}$. A model³ of $\text{xnf}(\varphi)^P$ is $\sigma = \{\text{X}(b), d, \text{X}_w(a)\}$. \triangleleft

Let σ be a model of $\text{xnf}(\varphi)^P$. If $\sigma \cap \mathcal{V}_X(\varphi) = \emptyset$, then we define $\sigma|_X = \text{end}$, where end is a special symbol not occurring in \mathcal{V} ; otherwise, we define:

³Models of the propositional formula $\text{xnf}(\varphi)^P$ are sets of propositional variables (or, equivalently, traces with one time instant).

OPERATORS	ARBITRARY FORMULAS		NEGATION-FREE FORMULAS				
	PLAIN	$\circ, *$	PLAIN	$+, +$	$+, \max$	$\max, +$	\max, \max
$T \subseteq \{X_w, R\}$	NP-c	NP-c	in P	NP-c	NP-c	in P	in P
$\{X\} \subseteq T \subseteq \{X_w, X\}$	NP-c	NP-c	in P	NP-c	NP-c	NP-c	in P
$\{U\} \subseteq T \subseteq \{X_w, U, R\}$	PSPACE-c	PSPACE-c	in P	PSPACE-c	PSPACE-c	in P	in P
$\{X, U\} \subseteq T$ or $\{X, R\} \subseteq T$	PSPACE-c	PSPACE-c	in P	PSPACE-c	PSPACE-c	PSPACE-c	in P

Figure 3: Summary of complexity results in Section 3 for weighted LTL_f —the column PLAIN reports the results that were already known to hold for LTL_f without weights (see [Fionda and Greco, 2018]).

replacing every occurrence of $\neg x$, for each $x \in \mathcal{V}$, with \bar{x} ; and (2) replacing every subformula φ_s with $\varphi_s \wedge \psi$. Consider then the valuation $v = \langle +, \max, w \rangle$ where $w(x) = w(\bar{x}) = 1$, for each $x \in \mathcal{V}$. It can be checked that φ is satisfiable if, and only if, $\bar{\varphi}$ admits a model π such that $v(\pi) \leq |\mathcal{V}|$. \square

4 System Design and Implementation

The notion of weighted transition graph and its properties (cf. Lemma 3.3) provided us with an immediate approach to implement a prototype for LTL_f over weighted traces, which we named WELTL. At a high level, it consists of two modules:

WELTL_{df}: The module builds in a depth-first manner the weighted transition graph, by checking for a model/path whose weight is below the desired threshold. Moreover, as we deal with monotone structures, it avoids exploring paths whose current weight exceeds that threshold.

WELTL_{bf}: The module uses a complementary breadth-first approach, which is better suited to be supported via a mechanism for implicit navigation, that we implemented in terms of a rewriting in *Answer Set Programming*.

In our experiments, we have found that WELTL_{df} is more effective to identify unsatisfiable instances, while WELTL_{bf} can be more practical to compute models of satisfiable formulas. With this observation to hand, we have coupled the two approaches. Since the design of WELTL_{df} poses no special challenges (indeed, we adapted the reasoner presented by [Li et al., 2020]), we next focus on the description of WELTL_{bf}, by subsequently presenting results of experimental activity.

4.1 WELTL_{bf} (and ASP in a Nutshell)

The architecture of WELTL_{bf} is depicted in Figure 4. Given a formula φ , it starts by computing a model π of minimum length, or immediately checks that the formula is not satisfiable—this task is delegated to the model checker made available by [Fionda and Greco, 2018]. If $v(\pi) \leq \tau$, then we are done. Otherwise, WELTL_{bf} starts its (implicit) breadth-first exploration by looking for a model whose length ℓ coincides with the length of π and such that the corresponding value is τ at most. This is carried out by leveraging a rewriting⁴ in terms of an Answer Set Program $P(\varphi, \ell, w)$.

For the sake of completeness, we recall that an ASP program [Brewka et al., 2011] is made of (a combination of): (i) facts of the form “*head*.”; (ii) normal rules of the form “*head* :- *body*.”; (iii) choice rules of the form “{*head*} :- *body*.”; (iv) constraints of the form “:- *body*.”, where *head* is an atom and *body* is a set of (possibly negated)




Figure 4: Conceptual architecture of WELTL_{bf}.

atoms, also including aggregate functions, such as *#sum* or *#max* [Faber et al., 2011]. Atoms can be made over variables (strings starting with uppercase letter) or constants (non-negative integers or strings starting with lowercase letters). The semantics is given in terms of its *answer sets*, that is, sets A of ground atoms such that: (i) *head* is in A ; (ii) whenever the *body* is true, *head* is in A ; (iii) *head* can be in A whenever the *body* is true; (iv) the *body* must be false.

The ASP rewriting $P(\varphi, \ell, \tau)$ uses the parse tree of φ : a fact $\text{var}(v_i)$ is included, for each node v and instant $i \in \{0, \dots, \ell - 1\}$. Then, a choice rule ($\{\text{selected}(X)\}:-\text{var}(X).$) is used to guess the subformulas of φ (represented by the nodes in the parse tree) that hold at i ; and, to check that such formulas actually hold, the definitions in Section 2 are straightforwardly translated in terms of constraints. Eventually, the ASP paradigm supports cost constraints that can encode the binary operators used in the valuation structures (in particular, they support “+” and “max”), so that the program can also enforce that answer set programs are associated with traces whose v -weight is τ at most. This is next exemplified.

Example 4.1 For an illustration of the rewriting, consider the formula $\varphi = a \wedge X(b)$, and the structure $v = \langle +, +, w \rangle$ where $w(a) = w(b) = 3$. Then, the ASP program $P(\varphi, 2, \tau)$ is shown in Figure 5. The parse tree has 4 nodes, denoted as *and*, *a*, *x*, and *b*. Then, rule (1) guesses the atoms corresponding to these nodes and to the time instants 0 and 1. Rule (2) forces that *selected*(*and*₀) holds, i.e., the whole formula holds at the initial time instant. Rules (3) and (4) define the semantics




Figure 5: ASP encoding of Example 4.1.

⁴Details on the encoding are in the Supplementary Material.

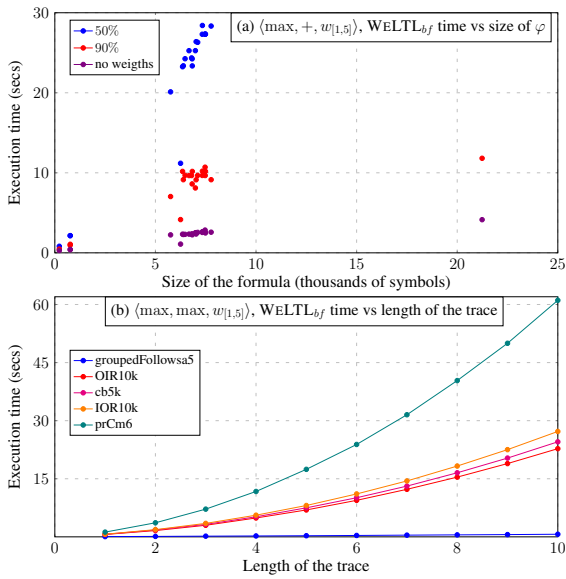


Figure 6: Execution times of over the benchmarks.

of the \wedge connective. Rule (5) defines the semantics of the X operator; in particular, note that when x_0 is selected, then b_1 must be selected too. Finally, rule (6) prevents that x_1 is selected, because we are considering $\ell = 2$ (and, hence, there is no subsequent time instant). The gray right box reports the rules encoding cost constraints. An atom weight ($\alpha_i, w(\alpha)$) is included, for each propositional variable α and time instant i . A cost constraint states that the sum of the weights must not exceeded τ . It is immediate to check that there is no answer set (hence, trace) satisfying the above rules for $\tau < 6$. \triangleleft

Finally, the ASP engine *clingo* [Gebser *et al.*, 2018] is executed on the rewriting $P(\varphi, \ell, w)$ (see Figure 4). If no answer set with the desired length ℓ and weight is found, then we explore the successive level. The process ends if either we actually find a solution, or if ℓ exceeds the theoretical length $\text{thB}(\varphi)$ of the longest path in the weighted transition graph, which is in general exponential w.r.t. $\|\varphi\|$, and polynomially bounded in some syntactic fragments (cf. Theorem 3.5).

4.2 Experimental Results

For the evaluation, we used a benchmark taken from the business process domain, namely a dataset of 109 satisfiable DECLARE formulas [van der Aalst *et al.*, 2009] obtained by mining all DECLARE constraints that hold on a number of logs made available by the IEEE Task-Force on Process Mining (<http://datacentrum.3tu.nl>). The DECLARE constraints are then translated into their corresponding LTL_f formulas [van der Aalst *et al.*, 2009]. For each formula, we randomly generated the variable weights in the intervals $[1, 5]$ and $[1, 10]$ and considered all combinations of \max and $+$ operators, thus ending up with a dataset of 872 weighted settings. Moreover, for each weighted setting, we considered four different cost bounds obtained as follows: (i) we compute an arbitrary model π ; (ii) we build a trace π' whose length n is the same as π and such that $\pi'_i = \mathcal{V}$ for each $i \in \{0, \dots, n-1\}$; (iii) we set the cost bounds to the 25%, 50%, 75% and 90% of $v(\pi')$. Clearly enough, for the reduced

$\langle o, *, w \rangle$	Solved formulas			
	25%	50%	75%	90%
$\langle \max, \max, w_{[1,5]} \rangle$	24 (0)	24 (0)	24 (0)	33 (14)
$\langle \max, \max, w_{[1,10]} \rangle$	24 (0)	24 (0)	25 (2)	42 (25)
$\langle \max, +, w_{[1,5]} \rangle$	25 (2)	96 (87)	109 (105)	109 (105)
$\langle \max, +, w_{[1,10]} \rangle$	26 (3)	104 (94)	109 (105)	109 (107)
$\langle +, \max, w_{[1,5]} \rangle$	109 (106)	109 (109)	109 (109)	109 (109)
$\langle +, \max, w_{[1,10]} \rangle$	109 (103)	109 (109)	109 (109)	109 (109)
$\langle +, +, w_{[1,5]} \rangle$	109 (108)	109 (109)	109 (109)	109 (109)
$\langle +, +, w_{[1,10]} \rangle$	109 (108)	109 (109)	109 (109)	109 (109)

Table 1: Summary of results. Columns “Solved formulas” report the number of formulas for which WELTL computes a solution. For the remaining formulas, WELTL terminates in timeout.

weight bounds, formulas quickly became unsatisfiable. Experiments have been carried out on a Intel CPU 2,4 GHz with 16GB RAM. Time and memory were limited to 600 seconds and 15GB, respectively. In particular, we first run WELTL_{bf} (with 300 seconds of timeout); and if no solution has been returned, then we run WELTL_{df} (again with the same timeout).

Table 1 reports a summary of the obtained results. For each valuation structure and each cost bound we report the number of formulas (out of 109) that are solved by WELTL within the time limit of 600 seconds. Each cell of the table reports inside the parenthesis the number of instances solved by WELTL_{bf} during the first 300 seconds. All the instances solved by WELTL_{bf} are satisfiable (i.e., a trace satisfying the cost bound has been found), while all the instances solved by WELTL_{df} during the last 300 seconds are unsatisfiable.

To get further insights on WELTL_{bf}, we report in Figure 6 the running times obtained on some instances. In particular, Figure 6.(a) reports the time to solve some satisfiable formulas when no weights are considered (i.e., standard satisfiability) and according to the valuation structure $\langle \max, +, w_{[1,5]} \rangle$ and the weight bounds 50% and 90%, respectively. The 22 formulas reported in the chart are those for which the model found at 90% did not meet the weight requirement at 50% and required to search for additional models. It can be noted that the time grows with the size of the formula and by lowering the cost bound. Figure 6.(b) reports the time required by WELTL_{bf} to check the existence of models for $\langle \max, \max, w_{[1,5]} \rangle$ and the bound 25% for 5 formulas, and by fixing different lengths ℓ in the encoding. The theoretical exponential dependency is nicely handled by the ASP engine.

5 Conclusion

A framework to express LTL_f properties in quantitative settings has been proposed and studied, based on the concept of valuation structures as mechanisms to equip each trace with a weight. While our approach departs from most of the quantitative LTL (and LTL_f) extensions proposed in the literature, a perspective similar to ours has been taken previously by [Lahijanian *et al.*, 2015]—though in the specific context of temporal planning with partial satisfaction (and without providing any complexity analysis and system prototype). In particular, a clear picture of the complexity issues arising when reasoning over weighted finite traces has been depicted and an efficient reasoning engine has been made available. On the theoretical level, instead, it remains open to study LTL_f over valuation structures that are not monotone, by characterizing its computational properties and by defining and implementing computation algorithms.

References

- [Almagor *et al.*, 2016] Shaull Almagor, Udi Boker, and Orna Kupferman. Formally reasoning about quality. *Journal of the ACM*, 63(3):24:1–24:56, 2016.
- [Alur *et al.*, 2001] Rajeev Alur, Kousha Etessami, Salvatore La Torre, and Doron A. Peled. Parametric temporal logic for “model measuring”. *ACM Trans. on Computational Logics*, 2(3):388–407, 2001.
- [Baier *et al.*, 2014] Christel Baier, Joachim Klein, Sascha Klüppelholz, and Sascha Wunderlich. Weight monitoring with linear temporal logic: complexity and decidability. In *Proc. of CSL-LICS*, pages 11:1–11:10. ACM, 2014.
- [Bouyer *et al.*, 2014] Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In *Proc. of CONCUR*, pages 266–280, 2014.
- [Brewka *et al.*, 2011] Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczynski. Answer set programming at a glance. *Commun. ACM*, 54(12):92–103, 2011.
- [Camacho *et al.*, 2018b] Alberto Camacho, Meghyn Bienvenu, and Sheila A. McIlraith. Finite LTL synthesis with environment assumptions and quality measures. In *Proc. of KR*, pages 454–463, 2018.
- [de Alfaro *et al.*, 2004] Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching metrics for quantitative transition systems. In *Proc. of ICALP*, pages 97–109, 2004.
- [de Alfaro *et al.*, 2005] Luca de Alfaro, Marco Faella, Thomas A. Henzinger, Rupak Majumdar, and Mariëlle Stoelinga. Model checking discounted temporal properties. *Theor. Comput. Sci.*, 345(1):139–170, 2005.
- [De Giacomo and Vardi, 2013] Giuseppe De Giacomo and M. Y. Vardi. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In *Proc. of IJCAI*, pages 854–860.
- [De Giacomo *et al.*, 2014a] Giuseppe De Giacomo, Riccardo De Masellis, Marco Grasso, Fabrizio M. Maggi, and Marco Montali. Monitoring Business Metaconstraints Based on LTL and LDL for Finite Traces. In *Proc. of BPM*, pages 1–14, 2014.
- [De Giacomo *et al.*, 2014b] Giuseppe De Giacomo, Riccardi De Masellis, and Marco Montali. Reasoning on LTL on Finite Traces: Insensitivity to Infiniteness. In *Proc. of AAAI*, pages 1027–1033, 2014.
- [De Giacomo and Rubin, 2018] Giuseppe De Giacomo and Sasha Rubin. Automata-theoretic foundations of FOND planning for LTLf and LDLf goals. In *Proc. of IJCAI*, pages 4729–4735, 2018.
- [De Giacomo and Vardi, 2015] Giuseppe De Giacomo and Moshe Y. Vardi. Synthesis for LTL and LDL on finite traces. In *Proc. of IJCAI*, pages 1558–1564, 2015.
- [Droste and Rahonis, 2016] Manfred Droste and George Rahonis. Weighted linear dynamic logic. In *Proc. of EPTCS*, volume 226, pages 149–163, 2016.
- [Droste and Vogler, 2012] Manfred Droste and Heiko Vogler. Weighted automata and multi-valued logics over arbitrary bounded lattices. *TCS*, 418:14–36, 2012.
- [Faber *et al.*, 2011] Wolfgang Faber, Gerald Pfeifer, and Nicola Leone. Semantics and complexity of recursive aggregates in answer set programming. *Artif. Intell.*, 175(1):278–298, 2011.
- [Faella *et al.*, 2008] Marco Faella, Axel Legay, and Mariëlle Stoelinga. Model checking quantitative linear time logic. *E. N. in Theor. Comput. Sci.*, 220(3):61–77, 2008.
- [Fionda and Greco, 2018] Valeria Fionda and Gianluigi Greco. LTL on finite and process traces: Complexity results and a practical reasoner. *J. Artif. Intell. Res.*, 63:557–623, 2018.
- [Gebser *et al.*, 2018] Martin Gebser, Nicola Leone, Marco Maratea, Simona Perri, Francesco Ricca, and Torsten Schaub. Evaluation techniques and systems for answer set programming: a survey. In *Proc. of IJCAI’18*, pages 5450–5456, 2018.
- [Kuperberg, 2014] Denis Kuperberg. Linear temporal logic for regular cost functions. *Log. Methods Comput. Sci.*, 10(1), 2014.
- [Kupferman and Lustig, 2007] Orna Kupferman and Yoav Lustig. Lattice automata. In *Proc. of VMCAI*, pages 199–213, 2007.
- [Kupferman *et al.*, 2009] Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. From liveness to promptness. *Formal Methods in System Design*, 34(2):83–103, 2009.
- [Lahijanian *et al.*, 2015] Morteza Lahijanian, Shaull Almagor, Dror Fried, Lydia E. Kavragi, and Moshe Y. Vardi. This time the robot settles for a cost: A quantitative approach to temporal logic planning with partial satisfaction. In *Proc. of AAAI*, pages 3664–3671, 2015.
- [Li *et al.*, 2020] Jianwen Li, Geguang Pu, Yueling Zhang, Moshe Y. Vardi, and Kristin Y. Rozier. Sat-based explicit ltl satisfiability checking. *Artif. Intell.*, 289:103369, 2020.
- [Mandralli, 2012] Eleni Mandrali. Weighted LTL with discounting. In *Proc. of CIAA’12*, pages 353–360, 2012.
- [Montali *et al.*, 2010] Marco Montali, Maja Pesic, Wil M. P. van der Aalst, Federico Chesani, Paola Mello, and Sergio Storari. Declarative specification and verification of service choreographiess. *ACM Trans. on the Web*, 4(1):3:1–3:62, 2010.
- [Pnueli, 1977] Amir Pnueli. The Temporal Logic of Programs. In *Proc. of FOCS*, pages 46–57, 1977.
- [van der Aalst *et al.*, 2009] W.M.P. van der Aalst, M. Pesic, and H. Schonenberg. Declarative Workflows: Balancing Between Flexibility and Support. *Computer Science - Research and Development*, 23(2):99–113, 2009.
- [Xiao *et al.*, 2021] Shengping Xiao, Jianwen Li, Shufang Zhu, Yingying Shi, Geguang Pu, and Moshe Y. Vardi. On-the-fly synthesis for LTL over finite traces. In *Proc. of AAAI*, pages 6530–6537, 2021.
- [Zhu *et al.*, 2020] Shufang Zhu, Giuseppe De Giacomo, Geguang Pu, and Moshe Y. Vardi. Ltlf synthesis with fairness and stability assumptions. In *Proc. of AAAI*, pages 3088–3095, 2020.