

Time window temporal logic [☆]Cristian-Ioan Vasile ^{a,*}, Derya Aksaray ^{b,2}, Calin Belta ^b^a Division of Systems Engineering, Boston University, Brookline, MA 02446, United States^b Department of Mechanical Engineering, Boston University, Boston, MA 02215, United States

ARTICLE INFO

Article history:

Received 12 March 2016

Received in revised form 19 June 2017

Accepted 17 July 2017

Available online 24 July 2017

Communicated by J.-F. Raskin

Keywords:

Timed temporal logic

Temporal relaxation

Controller synthesis

Verification

Finite state automata

Unambiguous languages

ABSTRACT

This paper introduces *time window temporal logic* (TWTL), a rich expressive language for describing various time bounded specifications. In particular, the syntax and semantics of TWTL enable the compact representation of serial tasks, which are prevalent in various applications including robotics, sensor systems, and manufacturing systems. This paper also discusses the relaxation of TWTL formulae with respect to the deadlines of the tasks. Efficient automata-based frameworks are presented to solve synthesis, verification and learning problems. The key ingredient to the presented solution is an algorithm to translate a TWTL formula to an annotated finite state automaton that encodes all possible temporal relaxations of the given formula. Some case studies are presented to illustrate the expressivity of the logic and the proposed algorithms.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Temporal logics provide mathematical formalisms to reason about (concurrent) events in terms of time. Due to their rich expressivity, they have been widely used as specification languages to describe properties related to correctness, termination, mutual exclusion, reachability, or liveness [34]. Recently, there has been great interest in using temporal logic formulae in the analysis and control of dynamical systems. For example, linear temporal logic (LTL) [5] has been extensively used in motion planning and control of robotic systems, e.g., [42,20,1,45,6,44,22,11,27,30].

In some real-world applications, the tasks may involve some time constraints (e.g., [38,36]). For example, consider a robot that is required to achieve the following tasks: every visit to *A* needs to be immediately followed by visiting *B* within 5 time units; two consecutive visits to *A* need to be at least 10 time units apart; or visiting *A* and visiting *B* need to be completed within 15 time units. Such tasks cannot be described by LTL formulae since LTL cannot deal with temporal properties with explicit time constraints. Therefore, bounded temporal logics are used to capture the time constraints over the tasks. Examples are bounded linear temporal logic (BLTL) [39,18], metric temporal logic (MTL) [26], and signal temporal logic (STL) [33].

[☆] This work was supported in part by NSF grants numbers NRI-1426907, NSF CMMI-1400167, and ONR grant number N00014-14-1-0554.

* Corresponding author.

E-mail addresses: cvasile@mit.edu (C.-I. Vasile), daksaray@mit.edu (D. Aksaray), cbelta@bu.edu (C. Belta).

¹ Present address: Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA.

² Present address: MIT Computer Science and Artificial Intelligence Laboratory, The Stata Center, Building 32, 32 Vassar Street, Cambridge, MA 02139, USA.

In this paper, we propose a specification language called *time window temporal logic* (TWTL). The semantics of TWTL is rich enough to express a wide variety of time-bounded specifications, e.g., “monitor A for 3 time units within the time interval $[0, 5]$ and after that monitor B for 2 time units within $[4, 9]$ ”. This logic was defined in our previous conference papers [43,2], and used to specify persistent surveillance tasks for multi-robot systems. Moreover, we define a notion of *temporal relaxation* of a TWTL formula, which is a quantity computed over the time intervals of a given TWTL formula. In this respect, if the temporal relaxation is: *negative*, then the tasks expressed in the formula should be completed before their designated time deadlines (i.e., satisfying the relaxed formula implies the satisfaction of a more strict formula than the original formula); *zero*, then the relaxed formula is exactly the same as the original formula; *positive*, then some tasks expressed in the formula are allowed to be completed after their original time deadlines (i.e., satisfying the relaxed formula may imply the violation of the original formula or the satisfaction of a less strict formula).

In this paper, we present an automata-based framework to solve verification, synthesis, and learning problems that involve TWTL specifications. One property of TWTL specifications we exploit in the proposed solutions is that the associated languages are finite. In the theoretical computer science literature, finite languages and the complexity of constructing their corresponding automata have been extensively studied [32,16,7,12,9]. One of the main benefits of the proposed framework is its capability to efficiently construct the annotated automata that can encode not only the original formula but also all temporal relaxations of the given formula. Such an efficient construction mainly stems from the proposed algorithms that are specifically developed for TWTL formulae.

The proposed language TWTL has several advantages over existing temporal logics. First, in many robotics missions, a desired specification can be represented in a more compact and comprehensible way in TWTL than BLTL, MTL, or STL. For example, deadlines expressed in a TWTL formula indicate the exact time bounds as opposed to an STL formula where the time bounds can be shifted. Consider a specification as “stay at A for 4 time steps within the time window $[0, 10]$ ”, which can be expressed in TWTL as $[H^4A]^{[0,10]}$. The same specification can be expressed in STL as $F_{[0,10-4]}G_{[0,4]}A$ where the outermost time window needs to be modified with respect to the inner time window. Furthermore, compared to BLTL and MTL, the existence of an explicit concatenation operator results in a more compact representation for serial tasks that are prevalent in various applications including robotics, sensor systems, and manufacturing systems. Under some mild assumptions, we provide a very efficient (linear-time) algorithm to handle concatenation of tasks. In general, the complexity associated with the concatenation operation is exponential in the worst case, even for finite languages [32].

Second, the notion of temporal relaxation enables a generic framework to construct the automaton of all possible relaxations of a TWTL formula. In literature, there are some studies investigating the control synthesis problems for minimal violations of LTL fragments [37,40,41,31,14]. In contrast to existing works, the annotated automaton proposed in this paper can encode all possible temporal relaxations of a given formula. Accordingly, such an automaton can be used in a variety of problems related to synthesis, verification, and learning to satisfy minimally relaxed formulae. Third, we show that the complexity of constructing the automata for a given TWTL formula is independent of the corresponding time bounds. To achieve this property, we exploit the structure of finite languages encoded by TWTL formulae.

We present a set of provably-correct algorithms to construct the automaton of a given TWTL formula (both for the relaxed and unrelaxed cases). We formulate a generic problem in terms of temporal relaxation of a TWTL formula, which can be specialized into problems such as verification, synthesis, and learning. We developed a Python package to solve these three problems, which is available for download from hyess.bu.edu/twtl.

2. Preliminaries

In this section, we introduce the notation and briefly review the main concepts from formal languages, automata theory, and formal verification. For a detailed exposition of these topics, the reader is referred to [5,17] and the references therein.

Given $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$, $n \geq 2$, the relationship $\mathbf{x} \sim \mathbf{x}'$, where $\sim \in \{<, \leq, >, \geq\}$, is true if it holds pairwise for all components. $\mathbf{x} \sim a$ denotes $\mathbf{x} \sim a\mathbf{1}_n$, where $a \in \mathbb{R}$ and $\mathbf{1}_n$ is the n -dimensional vector of all ones. The extended set of real numbers is denoted by $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$.

Let Σ be a finite set. We denote the cardinality and the power set of Σ by $|\Sigma|$ and 2^Σ , respectively. A *word* over Σ is a finite or infinite sequence of elements from Σ . In this context, Σ is also called an *alphabet*. The length of a word w is denoted by $|w|$ (e.g., $|w| = \infty$ if w is an infinite word). Let $k, i \leq j$ be non-negative integers. The k -th element of w is denoted by w_k , and the sub-word w_i, \dots, w_j is denoted by $w_{i,j}$. A set of words over an alphabet Σ is called a *language* over Σ . The languages of all finite and infinite words over Σ are denoted by Σ^* and Σ^ω , respectively.

Definition 2.1 (*Prefix language*). Let \mathcal{L}_1 and \mathcal{L}_2 be two languages. We say that \mathcal{L}_1 is a prefix language of \mathcal{L}_2 if and only if every word in \mathcal{L}_1 is a prefix of some word in \mathcal{L}_2 , i.e., for each word $w \in \mathcal{L}_1$ there exists $w' \in \mathcal{L}_2$ such that $w = w'_{0,i}$, where $0 \leq i < |w'|$. The maximal prefix language of a language \mathcal{L} is denoted by $P(\mathcal{L}) = \{w_{0,i} \mid w \in \mathcal{L}, i \in \{0, \dots, |w| - 1\}\}$.

Definition 2.2 (*Unambiguous language*). A language \mathcal{L} is called unambiguous language if no proper subset L of \mathcal{L} is a prefix language of $\mathcal{L} \setminus L$.

The above definition immediately implies that a word in an unambiguous language can not be the prefix of another word. Moreover, it is easy to show that the converse is also true.

Definition 2.3 (*Language concatenation*). Let \mathcal{L}_1 be a language over finite words, and let \mathcal{L}_2 be a language over finite or infinite words. The concatenation language $\mathcal{L}_1 \cdot \mathcal{L}_2$ is defined as the set of all words ww' , where $w \in \mathcal{L}_1$ and $w' \in \mathcal{L}_2$.

Definition 2.4 (*Deterministic finite state automaton*). A deterministic finite state automaton (DFA) is a tuple $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta_{\mathcal{A}}, F_{\mathcal{A}})$, where:

- $S_{\mathcal{A}}$ is a finite set of states;
- $s_0 \in S_{\mathcal{A}}$ is the initial state;
- Σ is the input alphabet;
- $\delta_{\mathcal{A}} : S_{\mathcal{A}} \times \Sigma \rightarrow S_{\mathcal{A}}$ is the transition function;
- $F_{\mathcal{A}} \subseteq S_{\mathcal{A}}$ is the set of accepting states.

A transition $s' = \delta_{\mathcal{A}}(s, \sigma)$ is also denoted by $s \xrightarrow{\sigma}_{\mathcal{A}} s'$. A trajectory of the DFA $\mathbf{s} = s_0s_1 \dots s_{n+1}$ is generated by a finite sequence of symbols $\sigma = \sigma_0\sigma_1 \dots \sigma_n$ if $s_0 \in S_{\mathcal{A}}$ is the initial state of \mathcal{A} and $s_k \xrightarrow{\sigma_k}_{\mathcal{A}} s_{k+1}$ for all $k \geq 0$. The trajectory generated by σ is also denoted by $s_0 \xrightarrow{\sigma}_{\mathcal{A}} s_{n+1}$. A finite input word σ over Σ is said to be accepted by a finite state automaton \mathcal{A} if the trajectory of \mathcal{A} generated by σ ends in a state belonging to the set of accepting states, i.e., $F_{\mathcal{A}}$. A DFA is called *blocking* if the $\delta_{\mathcal{A}}(s, \sigma)$ is a partial function, i.e., the value of the function is not defined for all values in the domain. A blocking automaton rejects words σ if there exists $k \geq 0$ such that $s_k \xrightarrow{\sigma_k}_{\mathcal{A}} s_{k+1}$ is not defined. The (*accepted*) *language* corresponding to a DFA \mathcal{A} is the set of accepted input words, which we denote by $\mathcal{L}(\mathcal{A})$.

Definition 2.5 (*Transition system*). A transition system (TS) is a tuple $\mathcal{T} = (X, x_0, \Delta, AP, h)$, where:

- X is a finite set of states;
- $x_0 \in X$ is the initial state;
- $\Delta \subseteq X \times X$ is a set of transitions;
- AP is a set of properties (atomic propositions);
- $h : X \rightarrow 2^{AP}$ is a labeling function.

We also denote a transition $(x, x') \in \Delta$ by $x \rightarrow_{\mathcal{T}} x'$. A *trajectory* (or run) of the system is an infinite sequence of states $\mathbf{x} = x_0x_1 \dots$ such that $x_k \rightarrow_{\mathcal{T}} x_{k+1}$ for all $k \geq 0$. A state trajectory \mathbf{x} generates an *output trajectory* $\mathbf{o} = o_0o_1 \dots$, where $o_k = h(x_k)$ for all $k \geq 0$. The (*generated*) *language* corresponding to a TS \mathcal{T} is the set of all generated output words, which we denote by $\mathcal{L}(\mathcal{T})$.

3. Time window temporal logic

Time window temporal logic (TWTL) was first introduced in the conference paper [43] as a rich specification language for robotics applications. Besides robotics, TWTL can be used in various domains (e.g., manufacturing, control, software development) that involve specifications with explicit time bounds. In particular, TWTL formulae can express tasks, their durations, and their time windows. TWTL is a linear-time logic encoding sets of discrete-time sequences with values in a finite alphabet.

A TWTL formula is defined over a set of atomic propositions AP and has the following syntax:

$$\phi ::= H^d s \mid H^d \neg s \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg \phi_1 \mid \phi_1 \cdot \phi_2 \mid [\phi_1]^{[a,b]}$$

where s is either the “true” constant \top or an atomic proposition in AP ; \wedge , \vee , and \neg are the conjunction, disjunction, and negation Boolean operators, respectively; \cdot is the concatenation operator; H^d with $d \in \mathbb{Z}_{\geq 0}$ is the *hold* operator; and $[]^{[a,b]}$ is the *within* operator, $a, b \in \mathbb{Z}_{\geq 0}$ and $a \leq b$.

The semantics of the operators is defined with respect to the finite subsequences of a (possibly infinite) word \mathbf{o} over 2^{AP} . Let \mathbf{o}_{t_1, t_2} be the subsequence of \mathbf{o} , which starts at time $t_1 \geq 0$ and ends at time $t_2 \geq t_1$. The *hold* operator $H^d s$ specifies that $s \in AP$ should be repeated for d time units. The semantics of $H^d \neg s$ is defined similarly, but for d time units only symbols from $AP \setminus \{s\}$ should appear. For convenience, if $d = 0$ we simply write s and $\neg s$ instead of $H^0 s$ and $H^0 \neg s$, respectively. The word \mathbf{o}_{t_1, t_2} satisfies $\phi_1 \wedge \phi_2$, $\phi_1 \vee \phi_2$, or $\neg \phi$ if \mathbf{o}_{t_1, t_2} satisfies both formulae, at least one formula, or does not satisfy the formula, respectively. The *within* operator $[\phi]^{[a,b]}$ bounds the satisfaction of ϕ to the time window $[a, b]$. The concatenation operator $\phi_1 \cdot \phi_2$ specifies that first ϕ_1 must be satisfied, and then immediately ϕ_2 must be satisfied.

Formally, the semantics of TWTL formulae is defined recursively as follows:

$$\begin{aligned} \mathbf{o}_{t_1, t_2} \models H^d s & \quad \text{iff } s \in o_t, \forall t \in \{t_1, \dots, t_1 + d\} \wedge (t_2 - t_1 \geq d) \\ \mathbf{o}_{t_1, t_2} \models H^d \neg s & \quad \text{iff } s \notin o_t, \forall t \in \{t_1, \dots, t_1 + d\} \wedge (t_2 - t_1 \geq d) \\ \mathbf{o}_{t_1, t_2} \models \phi_1 \wedge \phi_2 & \quad \text{iff } (\mathbf{o}_{t_1, t_2} \models \phi_1) \wedge (\mathbf{o}_{t_1, t_2} \models \phi_2) \\ \mathbf{o}_{t_1, t_2} \models \phi_1 \vee \phi_2 & \quad \text{iff } (\mathbf{o}_{t_1, t_2} \models \phi_1) \vee (\mathbf{o}_{t_1, t_2} \models \phi_2) \end{aligned}$$

$$\begin{aligned}
\mathbf{o}_{t_1, t_2} &\models \neg\phi && \text{iff } \neg(\mathbf{o}_{t_1, t_2} \models \phi) \\
\mathbf{o}_{t_1, t_2} &\models \phi_1 \cdot \phi_2 && \text{iff } (\exists t = \arg \min_{t_1 \leq t < t_2} \{\mathbf{o}_{t_1, t} \models \phi_1\}) \wedge \\
&&& (\mathbf{o}_{t+1, t_2} \models \phi_2) \\
\mathbf{o}_{t_1, t_2} &\models [\phi]^{[a, b]} && \text{iff } \exists t \geq t_1 + a \text{ s.t. } \mathbf{o}_{t, t_1+b} \models \phi \wedge (t_2 - t_1 \geq b)
\end{aligned}$$

A word \mathbf{o} is said to satisfy a formula ϕ if and only if there exists $T \in \{0, \dots, |\mathbf{o}|\}$ such that $\mathbf{o}_{0, T} \models \phi$.

A TWTL formula ϕ can be verified with respect to a bounded word. Accordingly, we define the *time bound* of ϕ , i.e., $\|\phi\|$, as the maximum time needed to satisfy ϕ , which can be recursively computed as follows:

$$\|\phi\| = \begin{cases} \max(\|\phi_1\|, \|\phi_2\|) & \text{if } \phi \in \{\phi_1 \wedge \phi_2, \phi_1 \vee \phi_2\} \\ \|\phi_1\| & \text{if } \phi = \neg\phi_1 \\ \|\phi_1\| + \|\phi_2\| + 1 & \text{if } \phi = \phi_1 \cdot \phi_2 \\ d & \text{if } \phi \in \{H^d s, H^d \neg s\} \\ b & \text{if } \phi = [\phi_1]^{[a, b]} \end{cases} \quad (1)$$

We denote the language of all words satisfying ϕ by $\mathcal{L}(\phi)$. Note that TWTL formulae are used to specify prefix languages of either Σ^* or Σ^ω , where $\Sigma = 2^{AP}$. Moreover, the number of operators in a TWTL formula ϕ is denoted by $|\phi|$.

Note that the semantics of the concatenation operation requires that the right operand formula starts one time unit after the left formula is done, i.e., its satisfaction has been established. The semantics of the sequential satisfaction is captured in the semantics by the *argmin* constraint in the definition above. The importance of the constraint becomes apparent when time widows (*within* operators) are considered, where the moment of switching between tasks is not a priori fixed in the formula. To illustrate this, consider the simple formula $[A]^{[0, 1]} \cdot B$ over the set of atomic propositions $\{A, B\}$. The language of the formula is

$$\begin{aligned}
&\{(\{A\})(\{B\}), (\{A, B\})(\{B\}), (\{A\})(\{A, B\}), (\{A, B\})(\{A, B\}), (\emptyset, \{A\})(\{B\}), (\emptyset, \{A, B\})(\{B\}), (\emptyset, \{A\})(\{A, B\}), \\
&(\emptyset, \{A, B\})(\{A, B\}), (\{B\}, \{A\})(\{B\}), (\{B\}, \{A, B\})(\{B\}), (\{B\}, \{A\})(\{A, B\}), (\{B\}, \{A, B\})(\{A, B\})\},
\end{aligned}$$

where the parentheses delimit the parts of the words satisfying the left and right formulae of the concatenation.

Some examples of TWTL formulae for a robot servicing at some regions can be as follows:

– *servicing within a deadline*: “service A for 2 time units before 10”,

$$\phi_1 = [H^2 A]^{[0, 10]} \text{ and } \|\phi_1\| = 10. \quad (2)$$

– *servicing within time windows*: “service A for 4 time units within $[3, 8]$ and B for 2 time units within $[4, 7]$ ”,

$$\phi_2 = [H^4 A]^{[3, 8]} \wedge [H^2 B]^{[4, 7]} \text{ and } \|\phi_2\| = 8. \quad (3)$$

– *servicing in sequence*: “service A for 3 time units within $[0, 5]$ and after this service B for 2 time units within $[4, 9]$ ”,

$$\phi_3 = [H^3 A]^{[0, 5]} \cdot [H^2 B]^{[4, 9]} \text{ and } \|\phi_3\| = 15. \quad (4)$$

– *servicing in strict sequence*: “service A for 3 time units within $[0, 5]$ and after this service B for 2 time units within $[0, 3]$ and then C for 4 time units within $[0, 6]$ and do not service B and C before A , and C before B within $[0, 16]$ ”,

$$\phi_4 = [H^3 A]^{[0, 5]} \cdot [H^2 B]^{[0, 3]} \cdot [H^4 C]^{[0, 6]} \wedge \neg[B \cdot [A]^{[0, 15]} \vee C \cdot [A \vee B]^{[0, 15]}]^{[0, 16]} \text{ and } \|\phi_4\| = 16. \quad (5)$$

– *enabling conditions*: “if A is serviced for 2 time units within 9 time units, then B should be serviced for 3 time units within the same time interval (i.e., within 9 time units)”,

$$\phi_5 = [H^2 A \Rightarrow [H^3 B]^{[2, 5]}]^{[0, 9]} \text{ and } \|\phi_5\| = 9, \quad (6)$$

where \Rightarrow denotes implication.

In order to describe rich specifications, a temporal logic can be selected based on the expressivity of the logic and the complexity of the corresponding algorithms (e.g., for automata construction). In general, expressivity and complexity are coupled terms such that a logic with very rich expressivity has very high complexity. Furthermore, the easiness to express the specifications and to comprehend the meaning of the formulae is also a crucial aspect when choosing temporal logics. TWTL induces finite languages, and it has the same expressivity of BLTL.³ On the other hand, STL and MTL are more expressive languages than TWTL since they are developed for real-time systems and can express continuous-time properties.

³ The next operator (\mathbf{X}) is usually part of the syntax of BLTL, but may be defined as $\mathbf{X}\phi \equiv \mathbf{G}^{\leq 1}\phi \vee (\neg\phi \wedge \mathbf{F}^{\leq 1}\phi)$. The nested next operator is denoted by \mathbf{X}^d , $d \in \mathbb{Z}_{\geq 0}$.

Table 1

The representation of (3) in TWTL, BLTL, and MTL.

TWTL	$[H^4 A]^{[3,8]} \wedge [H^2 B]^{[4,7]}$
BLTL	$X^3 F^{\leq 8-4-3} G^{\leq 4} A \wedge X^4 F^{\leq 7-2-4} G^{\leq 2} B$
MTL	$\bigvee_{i=3}^{8-4} G_{[i,i+4]} A \wedge \bigvee_{i=4}^{7-2} G_{[i,i+2]} B$

Table 2

The representation of (4) in TWTL, BLTL, and MTL.

TWTL	$[H^3 A]^{[0,5]} \cdot [H^2 B]^{[4,9]}$
BLTL	$F^{\leq 5-3} (G^{\leq 3} A \wedge X^{3+4} F^{\leq 9-2-4} G^{\leq 2} B)$
MTL	$\bigvee_{i=0}^{5-3} (G_{[i,i+3]} A \wedge \bigvee_{j=i+3+4}^{i+3+9-2} G_{[j,j+2]} B)$

Table 3

The representation of (5) in TWTL, and BLTL. Equivalent MTL formulae are too long, and thus were omitted.

TWTL	$[H^3 A]^{[0,5]} \cdot [H^2 B]^{[0,3]} \cdot [H^4 C]^{[0,6]} \wedge \neg[B \cdot A]^{[0,15]} \vee C \cdot [A \vee B]^{[0,15]}]^{[0,16]}$
BLTL	$(G^{\leq 3} (\neg(B \vee C))) \mathcal{U}^{\leq 5-3} A \wedge F^{\leq 5-3} (G^{\leq 3} A \wedge X^3 ((G^{\leq 2} \neg C) \mathcal{U}^{\leq 3-2} B \wedge F^{\leq 3-2} (G^{\leq 2} \wedge X^2 F^{\leq 6-4} G^{\leq 4} C)))$

TWTL provides some benefits over other time-bounded temporal logics. From the perspective of easiness to express specifications and to comprehend formulae, a main benefit of TWTL is the existence of concatenation, within, and hold operators. In particular, these operators lead to compact (shorter length) representation of specifications, which greatly improves the readability of the formulae. For example, consider the specifications in (3), (4) and (5), which are expressed in various temporal logics in Table 1, 2 and 3. Note that the TWTL formulae are short and comprehensible whereas an expert in formal methods might be required to create the other formulae to take into account the nested temporal operators, the shifted time windows, and the disjunction of numerous sub-formulae.

From the perspective of complexity, a main benefit of TWTL is the existence of explicit concatenation operator. In particular, the concatenation of two tasks can be expressed in other logics in a more sophisticated way than TWTL. In Table 2, we illustrate that the MTL formula contains a set of recursively defined sub-formulae connected by disjunctions whereas the BLTL formula contains nested temporal operators with conjunction. In both cases, dealing with the disjunction of numerous sub-formulae and the nested temporal operators with conjunction significantly increases the complexity of constructing the automaton (i.e., in exponential and quadratic ways, respectively [32]). On the other hand, we provide a linear-time algorithm in Sec. 7 to handle the concatenations of tasks under some mild assumptions.

Moreover, the automata construction algorithms in Sec. 7 are specifically developed for TWTL. Thus, an automaton for the satisfying language of a TWTL formula can be constructed directly (without translating it to another logic to use an off-the-shelf tool). For example, the authors of [39] translate a BLTL formula to a syntactically co-safe linear temporal logic (sclTL) formula [28] to use the automata construction tool *scheck* [29], which increases the complexity due to additional operations. Finally, for a given TWTL formula ϕ , we show that all possible temporally relaxed ϕ can be encoded to a very compact representation, which is enabled from the definition of temporal relaxation introduced in the next section.

4. Temporal relaxation

In this section, we introduce a *temporal relaxation* of a TWTL formula. This notion is used in Sec. 5 to formulate an optimization problem over temporal relaxations.

To illustrate the concept of temporal relaxation, consider the following TWTL formula:

$$\phi_1 = [H^1 A]^{[0,2]} \cdot [H^3 B \wedge [H^2 C]^{[0,4]}]^{[1,8]}. \quad (7)$$

In cases where ϕ_1 cannot be satisfied, one question is: what is the “closest” achievable formula that can be satisfied? Hence, we investigate relaxed versions of ϕ_1 . One way to do this is to relax the deadlines for the time windows, which are captured by the *within* operator. Accordingly, a relaxed version of ϕ_1 can be written as

$$\phi_1(\tau) = [H^1 A]^{[0,(2+\tau_1)]} \cdot [H^3 B \wedge [H^2 C]^{[0,(4+\tau_2)]}]^{[1,(8+\tau_3)]}, \quad (8)$$

where $\tau = (\tau_1, \tau_2, \tau_3) \in \mathbb{Z}^3$. Note that a critical aspect while relaxing the time windows is to preserve the feasibility of the formula. This means that all sub-formulae of ϕ enclosed by the *within* operators must take less time to satisfy than their corresponding time window durations.

Definition 4.1 (*Feasible TWTL formula*). A TWTL formula ϕ is called feasible, if the time window corresponding to each *within* operator is greater than the duration of the corresponding enclosed task (expressed via the *hold* operators).

Remark 4.1. Consider the formula in Eq. (8). For $\phi_1(\tau)$ to be a feasible TWTL formula, the following constraint must hold: (i) $2 + \tau_1 \geq 1$; (ii) $4 + \tau_2 \geq 2$ and (iii) $7 + \tau_3 \geq \max\{3, 4 + \tau_2\}$. Note that τ may be non-positive. In such cases, $\phi_1(\tau)$ becomes a stronger specification than ϕ_1 , which implies that the sub-tasks are performed ahead of their actual deadlines.

Let ϕ be a TWTL formula. Then, a τ -relaxation of ϕ is defined as follows:

Definition 4.2 (τ -Relaxation of ϕ). Let $\tau \in \mathbb{Z}^m$, where m is the number of *within* operators contained in ϕ . The τ -relaxation of ϕ is a feasible TWTL formula $\phi(\tau)$, where each subformula of the form $[\phi_i]^{[a_i, b_i]}$ is replaced by $[\phi_i]^{[a_i, b_i + \tau_i]}$.

Remark 4.2. For any ϕ , $\phi(\mathbf{0}) = \phi$.

Definition 4.3 (Temporal relaxation). Given ϕ , let $\phi(\tau)$ be a feasible relaxed formula. The temporal relaxation of $\phi(\tau)$ is defined as $|\tau|_{TR} = \max_j(\tau_j)$.

Remark 4.3. If a word $o \models \phi(\tau)$ with $|\tau|_{TR} \leq 0$, then $o \models \phi$.

5. Optimization over temporal relaxation

In this section, first, we propose a generic optimization problem over temporal relaxations of a TWTL formula. Then, we show how this setup can be used to formulate verification, synthesis, and learning problems.

The objective of the following optimization problem is to find a feasible relaxed version of a TWTL formula that optimizes a cost function penalizing the sets of satisfying and unsatisfying words, and the vector of relaxations.

Problem 5.1. Let ϕ be a TWTL formula over the set of atomic propositions AP , and let \mathcal{L}_1 and \mathcal{L}_2 be any two languages over the alphabet $\Sigma = 2^{AP}$. Consider a cost function $F : \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0} \times \mathbb{Z}^m \rightarrow \mathbb{R}$, where m is the number of *within* operators contained in ϕ . Find τ such that $F(|\mathcal{L}(\phi(\tau)) \cap \mathcal{L}_1|, |\mathcal{L}(\neg\phi(\tau)) \cap \mathcal{L}_2|, \tau)$ is minimized.

The sets \mathcal{L}_1 and \mathcal{L}_2 are used to impose constraints on the sets of words satisfying and violating the relaxed formula $\phi(\tau)$.

5.1. Verification, synthesis, and learning

In the following, we formulate three specific problems related to verification, synthesis, and learning based on [Problem 5.1](#). The synthesis problem addressed in this paper follows a recent trend of methods that return policies with reasonable performance even in the case when the specification cannot be met. In literature, some synthesis problems are framed as an optimization problem where the objective is to find a solution satisfying the minimal relaxation of a given specification [\[37,40,41,21,14\]](#). Alternatively, some studies impose a hierarchical structure on the input specification based on some given priorities [\[37,40,14\]](#). As such, lower priority properties may be disregarded in case the original specification cannot be satisfied. Yet another approach is presented in [\[31\]](#) where the authors consider a desired specification and a method to “locally” mend solution strategies in case these become infeasible. Thus, the method avoids global re-synthesis. Note that in these approaches it is very hard to translate and evaluate relaxed policies with respect to the original specifications.

The objective of the synthesis problem formulated in this section is to find a control policy (or strategy) that results in the satisfaction of the original formula or its minimal relaxation in case of infeasibility. Our solution approach differs from existing studies [\[37,40,41,21,14\]](#) in that the relaxation is defined at a semantic level, i.e., the TWTL formulae are parametrized. The main benefit of our approach is that the results of a synthesis algorithm can be interpreted in the same semantics as the original specification without using an additional representation (e.g., automata) for the relaxed formulae.

The verification problem addressed in this paper checks if a systems satisfies the structure of a specification without considering the time parameters, i.e., the deadlines of the *within* operators. This formulation differs from the generic ones that consider properties with fixed (temporal or spatial) parameters. Verification problems involving parametric formulae were also considered in [\[46\]](#) for STL and in [\[3\]](#) for LTL properties. In [\[46\]](#), the authors consider a (dense-time) STL specification with a single parameter and the problem of estimating bounds for that parameter. The solution is obtained using an optimization procedure that is defined in terms of robustness degree for STL properties. The problems explored in [\[3\]](#) are closer to the ones proposed in this paper. However, both bounded and unbounded properties are considered in [\[3\]](#) and the focus of the exposition is geared towards establishing decidability and complexity bounds.

Lastly, we address a parameter learning problem where the goal is to learn the time parameters of a TWTL formula from a given data set. The parameter synthesis for PSTL formulae is tackled in [\[4,19\]](#). Moreover, *Temporal Logic Inference*, which is the problem of learning both the structure and parameters of properties, is considered in [\[23,13\]](#). In this paper, we focus only on the inference of deadlines for TWTL formulae from labeled data such that the misclassification rate is minimized.

5.1.1. Verification⁴

Given a transition system \mathcal{T} and a TWTL formula ϕ , we want to check if there exists a relaxed formula $\phi(\tau)$ such that all output words generated by \mathcal{T} satisfy $\phi(\tau)$.

In [Problem 5.1](#), we can set $\mathcal{L}_1 = \emptyset$ and $\mathcal{L}_2 = \mathcal{L}(\mathcal{T})$, and we choose the following cost function:

$$F(x, y, \tau) = 1 - \delta(y), \quad (9)$$

where $x, y \in \mathbb{Z}_{\geq 0}$ and $\delta(x) = \begin{cases} 1 & x = 0 \\ 0 & x \neq 0 \end{cases}$. The cost function in Eq. (9) has a single global minimum value at 0 which corresponds to the case $\mathcal{L}(\mathcal{T}) \cap \mathcal{L}(\neg\phi(\tau)) = \emptyset$.

5.1.2. Synthesis

Given a transition system \mathcal{T} and a TWTL formula ϕ , we want to find a policy (a trajectory of \mathcal{T}) that produces an output word satisfying a relaxed version $\phi(\tau)$ of the specification with minimal temporal relaxation $|\tau|_{TR}$.

In [Problem 5.1](#), we can set $\mathcal{L}_1 = \mathcal{L}(\mathcal{T})$ and $\mathcal{L}_2 = \emptyset$, and we choose the following cost function:

$$F(x, y, \tau) = \begin{cases} |\tau|_{TR} & x > 0 \\ \infty & \text{otherwise} \end{cases}, \quad (10)$$

where $x, y \in \mathbb{Z}_{\geq 0}$. The cost function in Eq. (10) is minimized by an output word of \mathcal{T} , which satisfies the relaxed version of ϕ with minimum temporal relaxation, see [Definition 4.3](#).

5.1.3. Learning

Let ϕ be a TWTL formula and \mathcal{L}_p and \mathcal{L}_n be two finite sets of words labeled as positive and negative examples, respectively. We want to find a relaxed formula $\phi(\tau)$ such that the misclassification rate, i.e., $|\{w \in \mathcal{L}_p \mid w \not\models \phi(\tau)\}| + |\{w \in \mathcal{L}_n \mid w \models \phi(\tau)\}|$, is minimized.

This case can be mapped to the generic formulation by setting $\mathcal{L}_1 = \mathcal{L}_n$, $\mathcal{L}_2 = \mathcal{L}_p$ and choosing the cost function

$$F(x, y, \tau) = x + y, \quad (11)$$

which captures the misclassification rate, where $x, y \in \mathbb{Z}_{\geq 0}$.

5.2. Overview of the solution

We propose an automata-based approach to solve the verification, synthesis, and learning problems defined above. Specifically, the proposed algorithm constructs an annotated DFA \mathcal{A}_∞ , which captures all temporal relaxations of the given formula ϕ , i.e., $\mathcal{L}(\mathcal{A}_\infty) = \mathcal{L}(\phi(\infty))$ (see [Definition 6.3](#) for the definition of $\phi(\infty)$). Note that the algorithm can also be used to construct a (normal) DFA \mathcal{A} which accepts the satisfying language of ϕ , i.e., $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\phi)$. Using the resulting DFA \mathcal{A}_∞ , we proceed in [Sec. 8](#) to solve the synthesis and verification problems using a product automaton approach. For the synthesis problem, we propose a recursive algorithm that computes a satisfying path with minimum temporal relaxation. The learning problem is solved by inferring the minimum relaxation for each trajectory and then combining these relaxations to ensure minimum misclassification rate.

6. Properties of TWTL

In this section, we present properties of TWTL formulae, their temporal relaxations, and their accepted languages.

In this paper, languages are represented in three ways: as TWTL formulae, as automata, and as sets. As one might expect, there is a duality between some operators of TWTL and set operations, i.e., conjunction, disjunction, and concatenation correspond to intersection, union, and concatenation languages, respectively. Negation may be mapped to complementation with respect to the language of all bounded words, where the bound is given by the time bound of the negated formula.

Proposition 6.1. *The following properties hold*

$$(\phi_1 \cdot \phi_2) \cdot \phi_3 = \phi_1 \cdot (\phi_2 \cdot \phi_3) \quad (12)$$

$$\phi_1 \cdot (\phi_2 \vee \phi_3) = (\phi_1 \cdot \phi_2) \vee (\phi_1 \cdot \phi_3) \quad (13)$$

$$[\phi_1 \vee \phi_2]^{[a,b]} = [\phi_1]^{[a,b]} \vee [\phi_2]^{[a,b]} \quad (14)$$

$$\neg(H^d p) = [\neg p]^{[0,d]} \quad (15)$$

⁴ This problem is not a verification problem in the usual sense, but rather finding a formula that is satisfied by all runs of a system.

$$[\phi_1]^{[a_1, b_1]} = \bigvee_{i=a_1}^{b_1} (H^{i-1} \top \cdot \phi_1) \quad (16)$$

$$[\phi_1]^{[a_1, b_1]} = (H^{a_1-1} \top) \cdot [\phi_1]^{[0, b_1-a_1]} \quad (17)$$

$$(H^{d_1} p) \cdot (H^{d_2} p) = H^{d_1+d_2+1} p \quad (18)$$

$$[\phi_1]^{[a, b]} \Rightarrow [\phi_1]^{[a, b+\tau]} \quad (19)$$

where ϕ_1 , ϕ_2 , and ϕ_3 are TWTL formulae, $p \in \{s, \neg s\}$, $s \in AP \cup \{\top\}$, and $a, b, a_1, b_1, d, d_1, d_2, \tau \in \mathbb{Z}_{\geq 0}$ such that $a \leq b$ and $1 \leq a_1 \leq b_1$.

Proof. These follow directly from the semantics of TWTL formulae. \square

Definition 6.1 (Disjunction-Free Within form). Let ϕ be a TWTL formula. We say that ϕ is in Disjunction-Free Within (DFW) form if for all *within* operators contained in the formula the associated enclosed subformulae do not contain any disjunction operators.

An example of a TWTL formula in DFW form is $\phi_1 = [H^2 A]^{[0,9]} \vee [H^5 B]^{[0,9]}$, while a formula not in DFW form is $\phi_2 = [H^2 A \vee H^5 B]^{[0,9]}$. However, ϕ_1 and ϕ_2 are equivalent by Eq. (14) of Proposition 6.1. The next proposition formalizes this property.

Proposition 6.2. For any TWTL formula ϕ , if the negation operators are only in front of the atomic propositions, then ϕ can be written in the DFW form.

Proof. The result follows from the properties of distributivity of Boolean operators and Proposition 6.1, which can be applied iteratively to move all disjunction operators outside the *within* operators. \square

In the following, we define the notion of unambiguous concatenation, which enables tracking of progress for sequential specifications. Specifically, if the property holds, then an algorithm is able to decide at each moment if the first specification has finished while monitoring the satisfaction of two sequential specifications.

Definition 6.2. Let \mathcal{L}_1 and \mathcal{L}_2 be two languages. We say that the language $\mathcal{L}_1 \cdot \mathcal{L}_2$ is an *unambiguous concatenation* if each word in the resulting language can be split unambiguously, i.e., $(L_1, \mathcal{L}_1, \mathcal{L}_1 \cdot (P(\mathcal{L}_2) \setminus \{\epsilon\}))$ is a partition of $P(\mathcal{L}_1 \cdot \mathcal{L}_2)$, where $L_1 = \{w_{0,i} \mid w \in \mathcal{L}_1, i \in \{0, \dots, |w| - 2\}\}$ and $P(L)$ denotes the maximal prefix language of L .

The three sets of the partition from Definition 6.2 may be thought as indicating whether the first specification is in progress, the first specification has finished, and the second specification is in progress, respectively.

Proposition 6.3. Consider two languages \mathcal{L}_1 and \mathcal{L}_2 . The language $\mathcal{L}_1 \cdot \mathcal{L}_2$ is an unambiguous concatenation if and only if \mathcal{L}_1 is an unambiguous language.

Proof. Let $(L_1, \mathcal{L}_1, \mathcal{L}_1 \cdot (P(\mathcal{L}_2) \setminus \{\epsilon\}))$ be a partition of $P(\mathcal{L}_1 \cdot \mathcal{L}_2)$ and L be a proper subset of \mathcal{L}_1 . Assume that there exists $w \in L$ and $w' \in \mathcal{L}_1 \setminus L$ such that $w = w'_{0,i}$, for some $i \in \{0, \dots, |w'| - 1\}$. It follows that $w \in L_1$, because $w \neq w'$. However, this contradicts the fact that L_1 and \mathcal{L}_1 are disjoint.

Conversely, let \mathcal{L}_1 be unambiguous and consider a word $w \in P(\mathcal{L}_1 \cdot \mathcal{L}_2)$. Assume that $w \in L_1 \cap \mathcal{L}_1$. It follows that $\{w\}$ is a prefix language for $\mathcal{L}_1 \setminus \{w\}$, which contradicts with the hypothesis that \mathcal{L}_1 is unambiguous. Similarly, if we assume that there exists $w \in P(\mathcal{L}_1) \cap (\mathcal{L}_1 \cdot (P(\mathcal{L}_2) \setminus \{\epsilon\}))$, then there exists $w', w'' \in \mathcal{L}_1$ such that w' is a prefix of w , w is a prefix of w'' , and $|w'| < |w| \leq |w''|$. Thus, we arrive again at a contradiction with the unambiguity of \mathcal{L}_1 . Thus, the three sets form a partition of $P(\mathcal{L}_1 \cdot \mathcal{L}_2)$. \square

Proposition 6.4. Consider two unambiguous languages \mathcal{L}_1 and \mathcal{L}_2 . The language $\mathcal{L}_1 \cup \mathcal{L}_2$ is unambiguous if and only if $\mathcal{L}_1 \cap P(\mathcal{L}_2) = \mathcal{L}_2 \cap P(\mathcal{L}_1) = \emptyset$.

Proof. Let $\mathcal{L}_1 \cup \mathcal{L}_2$ be unambiguous. Assume that there exists $w \in \mathcal{L}_1 \cap P(\mathcal{L}_2)$. It follows that $w \in P(\mathcal{L}_2) \subseteq P(\mathcal{L}_1 \cup \mathcal{L}_2)$ and $w \in \mathcal{L}_1 \cup \mathcal{L}_2$, which implies a contradiction with the hypothesis that $\mathcal{L}_1 \cup \mathcal{L}_2$ is unambiguous. Thus, we obtain $\mathcal{L}_1 \cap P(\mathcal{L}_2) = \emptyset$. Similarly, it follows that $\mathcal{L}_2 \cap P(\mathcal{L}_1) = \emptyset$.

Conversely, let $\mathcal{L}_1 \cap P(\mathcal{L}_2) = \mathcal{L}_2 \cap P(\mathcal{L}_1) = \emptyset$. Let $L = L_1 \cup L_2 \subset \mathcal{L}_1 \cup \mathcal{L}_2$ such that $L \neq \emptyset$, $L_1 \subset \mathcal{L}_1$ and $L_2 \subset \mathcal{L}_2$. It follows that $L_1 \cap P(\mathcal{L}_1) = \emptyset$ and $L_1 \cap P(\mathcal{L}_2) = \emptyset$. Thus, $L_1 \cap (P(\mathcal{L}_1) \cup P(\mathcal{L}_2)) = L_1 \cap P(\mathcal{L}_1 \cup \mathcal{L}_2) = \emptyset$. Similarly, we obtain $L_2 \cap P(\mathcal{L}_1 \cup \mathcal{L}_2) = \emptyset$. Finally, it follows that $L \cap P(\mathcal{L}_1 \cup \mathcal{L}_2) = (L_1 \cup L_2) \cap P(\mathcal{L}_1 \cup \mathcal{L}_2) = \emptyset$, which shows that L cannot be a prefix language of $\mathcal{L}_1 \cup \mathcal{L}_2$. Therefore, $\mathcal{L}_1 \cup \mathcal{L}_2$ is unambiguous. \square

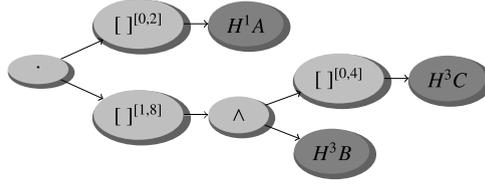


Fig. 1. An AST corresponding to the TWTL in Eq. (7). The intermediate light gray nodes correspond to the Boolean, concatenation, and *within* operators, while the dark gray leaf nodes represent the *hold* operators.

Proposition 6.5. *If \mathcal{L}_1 and \mathcal{L}_2 are unambiguous languages, then language $\mathcal{L}_1 \cap \mathcal{L}_2$ is unambiguous.*

Proof. Let $w \in \mathcal{L}_1 \cap \mathcal{L}_2$. It follows that $w \notin P(\mathcal{L}_1)$ and $w \notin P(\mathcal{L}_2)$. Therefore $w \notin P(\mathcal{L}_1 \cap \mathcal{L}_2)$, which implies $\mathcal{L}_1 \cap \mathcal{L}_2$ is unambiguous. \square

In the following results, we frequently use the notion of abstract syntax tree of a TWTL formula. An *Abstract Syntax Tree* (AST) of ϕ is denoted by $AST(\phi)$, where each leaf corresponds to a *hold* operator and each intermediate node corresponds to a Boolean, concatenation, or *within* operator. Given a TWTL formula ϕ , there might exist multiple AST trees that represent ϕ . In this paper, $AST(\phi)$ is assumed to be computed by an LL(*) parser [35]. The reader is referred to [17] for more details on AST and parsers. An example of an AST tree of Eq. (7) is illustrated in Fig. 1.

Proposition 6.6. *Let $\tau', \tau'' \in \mathbb{Z}^m$ such that $\phi(\tau')$ and $\phi(\tau'')$ are two feasible relaxed formulae, where m is the number of within operators in ϕ . If $\tau' \leq \tau''$, then $\phi(\tau') \Rightarrow \phi(\tau'')$.*

Proof. The proof follows by structural induction over $AST(\phi)$. The base case is trivial, since the leaves correspond to the *hold* operators. For the induction step, the result follows trivially if the intermediate node is associated with a Boolean or concatenation operator. The case of the *within* operator follows from Eq. (16) and (19) in Proposition 6.1. Let $\phi = [\phi_1]^{[a,b]}$. By the induction hypothesis, it holds $\phi_1(\tau'_1) \Rightarrow \phi_1(\tau''_1)$ for all $\tau'_1 \leq \tau''_1 \in \mathbb{Z}^{m-1}$. Let $\tau' \leq \tau'' \in \mathbb{Z}$, $\tau' = [\tau'^T \tau']^T$ and $\tau'' = [\tau''^T \tau'']^T$. Thus, $\phi(\tau') = \bigvee_{i=a}^{b+\tau'} (H^{i-1} \top \cdot \phi_1(\tau'_1)) \Rightarrow \bigvee_{i=a}^{b+\tau''} (H^{i-1} \top \cdot \phi_1(\tau''_1)) = [\phi_1(\tau''_1)]^{[a,b+\tau'']} \Rightarrow [\phi_1(\tau''_1)]^{[a,b+\tau'']}$, where the first implication follows from the induction hypothesis applied to each pair of delayed formulae, i.e., $(H^{i-1} \top \cdot \phi_1(\tau'_1)) \Rightarrow (H^{i-1} \top \cdot \phi_1(\tau''_1))$, for all $i \geq 0$. The second implication holds due to Eq. (19). \square

Definition 6.3. Given an output word \mathbf{o} , we say that \mathbf{o} satisfies $\phi(\infty)$, i.e., $\mathbf{o} \models \phi(\infty)$, if and only if $\exists \tau' < \infty$ s.t. $\mathbf{o} \models \phi(\tau')$.

The next corollary follows directly from Proposition 6.6.

Corollary 6.7. *Let $\tau < \infty$, then $\phi(\tau) \Rightarrow \phi(\infty)$, $\forall \tau$.*

Proposition 6.8. *Let $\phi(\tau')$ and $\phi(\tau'')$ be two feasible relaxed formulae. If $\tau' \leq \tau''$, then $\|\phi(\tau')\| \leq \|\phi(\tau'')\|$.*

Proof. The result follows by structural induction from Eq. (1) using a similar argument as in the proof of Proposition 6.6. \square

An important observation about TWTL is that the accepted languages corresponding to formulae are finite languages. In the following, we characterize such languages in terms of the associated automata.

Definition 6.4. A DFA is called *strict* if and only if (i) the DFA is blocking, (ii) all states reach a final state, and (ii) all states are reachable from the initial state.

Proposition 6.9. *Any DFA \mathcal{A} may be converted to a strict DFA in $O(|S_{\mathcal{A}}| \cdot |\Sigma|)$ time.*

Proof. States unreachable from the initial state can be identified by traversing the automaton graph from the initial state using either breath- or depth-first search. Similarly, the states not reaching a final state can be removed by traversing the automaton graph using the reverse direction of the transitions. Both operations take at most $O(|\delta_{\mathcal{A}}|) = O(|S_{\mathcal{A}}| \cdot |\Sigma|)$, since there are at most $|\Sigma|$ transitions outgoing from each state, where Σ is the alphabet of \mathcal{A} . \square

Note that a strict DFA is not necessarily minimal with respect to the number of states.

Proposition 6.10. *If \mathcal{L} is a finite language over an alphabet Σ , then the corresponding strict DFA is a directed acyclic graph (DAG). Moreover, given a (general) DFA \mathcal{A} , checking if its language $\mathcal{L}(\mathcal{A})$ is finite takes $O(|S_{\mathcal{A}}| \cdot |\Sigma|)$ time.*

Proof. For the first part, assume for the sake of contradiction that \mathcal{A} has a cycle. Then, we can form words in the accepted language by traversing the cycle $n \in \mathbb{Z}_{\geq 0}$ times before going to a final state. Note that the states in the cycle are reachable from the initial state and also reach a final state, because \mathcal{A} is a strict DFA. It follows that \mathcal{L} is infinite, which contradicts the hypothesis. Checking if a DFA \mathcal{A} is DAG takes $O(|S_{\mathcal{A}}| \cdot |\Sigma|)$ by using a topological sorting algorithm, because of the same argument as in [Proposition 6.9](#). \square

Corollary 6.11. *Let \mathcal{L} be a finite unambiguous language over the alphabet Σ and \mathcal{A} be its corresponding strict DFA. The following two statements hold:*

1. *if $s \in F_{\mathcal{A}}$, then the set of outgoing transitions of s is empty.*
2. *\mathcal{A} may be converted to a DFA with only one final states.*

Proof. Consider a final state $s \in F_{\mathcal{A}}$. Assume that there exists $s' \in S_{\mathcal{A}}$ such that $s \xrightarrow{\sigma}_{\mathcal{A}} s'$, where $\sigma \in \Sigma$. Since \mathcal{A} is strict, it follows that there is another final state $s'' \in F_{\mathcal{A}}$ which can be reached from s' . Next, we form the words w and w' leading to s and s'' passing through s' , respectively. Clearly, w is a prefix of w' , which implies that \mathcal{L} is not an unambiguous language. The second statement follows from the first by noting that in this case, merging all final states does not change the accepted language of the DFA \mathcal{A} . \square

The following result shows that any finite language can be expressed by either TWTL and BLTL. However, as shown in [Sec. 3](#), the two logics may not be equally concise.

Proposition 6.12. *For every finite language L , there exists a TWTL formula ϕ and a BLTL formula φ such that $L = \mathcal{L}(\phi) = \mathcal{L}(\varphi)$.*

Proof. First, note that both TWTL and BLTL have bounded time semantics, and therefore describe finite languages. Let AP be a finite set of atomic propositions and L a finite language over $\Sigma = 2^{AP}$. For every word $w = w_0, w_1, \dots \in L$, we can form the TWTL formula $\phi_w = \prod_{i=0}^{|w|-1} (\bigwedge_{\sigma \in w_i} \sigma)$, and BLTL formula $\varphi_w = \bigwedge_{i=0}^{|w|-1} (\mathbf{x}^i (\bigwedge_{\sigma \in w_i} \sigma))$, where \prod denotes iterated concatenation. Consider $\phi_L = \bigvee_{w \in L} \phi_w$ and $\varphi_L = \bigvee_{w \in L} \varphi_w$. It follows that $L = \mathcal{L}(\phi_w) = \mathcal{L}(\varphi_L)$. \square

7. Automata construction

In this section, we present a recursive procedure to construct DFAs for TWTL formulae and their temporal relaxations. The resulting DFA are used in [Sec. 8](#) to solve the proposed problems in [Sec. 5.1](#).

An important property of the automata representation proposed in this section is that it is independent of the time bounds (deadlines) in the formula. Thus, we obtain a parametric representation of all relaxed formulae from a given TWTL formula. Note that a similar (or parametric) representation may not be accommodated via SAT, SMT or MILP encoding schemes which depend on the exact time bounds of the encoded formulae.

Throughout the paper, a TWTL formula is assumed to have the following properties:

Assumption 1. Let ϕ be a TWTL. Assume that (i) negation operators are only in front of atomic propositions, and (ii) all sub-formulae of ϕ correspond to unambiguous languages.

The second part (ii) of [Assumption 1](#) is a desired property of specifications in practice, because it is related to the tracking of progress towards the satisfaction of the tasks. More specifically, if (ii) holds, then the end of each sub-formula can be determined unambiguously, i.e., without any look-ahead. Formulae that induce ambiguous languages correspond to bad specification. In [Sec. 7.6](#), we provide a method to check (ii) and associated complexity analysis.

As stated previously in [Sec. 3](#), TWTL formulae define prefix languages. Thus, the algorithms presented in this section compute automata with finite languages that terminate when their final states are reached.

7.1. Construction algorithm

In [\[43\]](#), a TWTL formula ϕ is translated to an equivalent sLTL formula, and then an off-the-shelf tool, such as *scheck* [\[29\]](#) and *spot* [\[10\]](#), is used to obtain the corresponding DFA. In this paper, we propose an alternative construction, shown in [Algorithm 1](#), with two main advantages: (i) the proposed algorithm is optimized for TWTL formulae so it is significantly faster than the method used in [\[43\]](#), and (ii) the same algorithm can be used to construct a special DFA, which captures all τ -relaxations of ϕ , i.e., the DFA \mathcal{A}_{∞} corresponding to $\phi(\infty)$.

[Algorithm 1](#) constructs the DFA recursively by traversing $AST(\phi)$ computed via an LL(*) parser [\[17,35\]](#) from the leaves to the root. If the parameter *inf* is true, then the returned DFA is an annotated DFA \mathcal{A}_{∞} corresponding to $\phi(\infty)$; otherwise a normal DFA \mathcal{A} is returned. Each operator has an associated algorithm ϱ_{\otimes} with $\otimes \in \{\wedge, \vee, \cdot, H, \infty, []\}$, which takes the DFAs corresponding to the operands (subtrees of the operator node in the AST) as input. Then, ϱ_{\otimes} returns the DFA that accepts

Algorithm 1: Translation algorithm – *translate*(·).

Input: ϕ – the specification as a TWTL formula in DFW form
Input: *inf* – flag specifying if the normal or annotated DFA is computed
Output: \mathcal{A} – translated DFA

```

1 if  $\phi = \phi_1 \otimes \phi_2$ , where  $\otimes \in \{\wedge, \vee, \cdot\}$  then
2   |  $\mathcal{A}_1 \leftarrow \text{translate}(\phi_1)$ ,  $\mathcal{A}_2 \leftarrow \text{translate}(\phi_2)$ 
3   |  $\mathcal{A} \leftarrow \mathcal{Q}_{\otimes}(\mathcal{A}_1, \mathcal{A}_2)$ 
4 else if  $\phi = H^d p$ , where  $p \in \{s, \neg s\}$  and  $s \in AP$  then
5   |  $\mathcal{A} \leftarrow \mathcal{Q}_H(p, d, AP)$ 
6 else if  $\phi = [\phi_1]^{[a,b]}$  then
7   |  $\mathcal{A}_1 \leftarrow \text{translate}(\phi_1)$ 
8   | if inf then  $\mathcal{A} \leftarrow \mathcal{Q}_{\infty}(\mathcal{A}_1, a, b)$ 
9   | else  $\mathcal{A} \leftarrow \mathcal{Q}_1(\mathcal{A}_1, a, b)$ 
10 return  $\mathcal{A}$ 

```

the formula associated with the operator node. In the following, we present all operators and related operations, such as annotating a DFA, relabeling the states of a DFA, or returning the truncated version of a DFA with respect to some given bound.

7.2. Annotation

The algorithms presented in this section use DFAs with some additional annotation. In this subsection, we introduce an annotated DFA and two algorithms, [Algorithm 3](#) and [Algorithm 2](#), that are used to (re)label DFAs and the associated annotation data, respectively.

We assume the following conventions to simplify the notation: (i) there is a global boolean variable *inf* accessible by all algorithms, which specifies whether the normal or the annotated DFAs are to be computed; (ii) in all algorithms, we have $\Sigma = 2^{AP}$; (iii) an element of $\Sigma \in \Sigma$ is called a *symbol* and is also a set of atomic propositions, $\sigma \subseteq AP$; (iv) a symbol σ is called *blocking* for a state s if there is no outgoing transition from s activated by σ .

7.2.1. Annotation

An *annotated DFA* is a tuple $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta, F_{\mathcal{A}}, T_{\mathcal{A}})$, where the first five components have the same meaning as in [Definition 2.4](#) and $T_{\mathcal{A}}$ is a tree that corresponds to the AST of the formula associated with the DFA. Each node T of the tree contains the following information:

1. $T.op$ is the operation corresponding to T ;
2. $T.I$ is the set of initial states of the automaton corresponding to T ;
3. $T.F$ is the set of final states of the automaton corresponding to T ;
4. $T.left$ and $T.right$ are the left and right child nodes of T , respectively.

Additionally, if $T.op$ is \vee (disjunction), then T has another attribute $T.choice$, which is explained in [Sec. 7.3.2](#).

Note that the associated trees are set to \emptyset and are ignored, if the normal DFAs are computed, i.e., *inf* is false.

The labels of the states change during the construction of the automata. [Algorithm 2](#) is used to update the labels stored in the data structures of the tree. The algorithm takes the tree T as input, a mapping m from the states to the new labels, and a boolean value e that specifies if the states are mapped to multiple new states. The first step is to convert the states' new labels to singleton sets if e is false (line 1). Then, the algorithm proceeds to process the tree recursively starting with T . The mapping m is then used to compute $t.I$ and $t.F$ by expanding each state to a set and then computing the union of the corresponding sets (lines 5–6). In the case of $op = \vee$, the three sets B , L , and R , which form the tuple $t.choices$ are also processed. The elements of all three sets are pairs of a state s and a symbol $\sigma \in \Sigma$. [Algorithm 2](#) converts the states of all these pairs in the tree sets (lines 7–12).

7.2.2. Relabeling a DFA

The [Algorithm 3](#) relabels the states of a DFA \mathcal{A} with labels given by the mapping m . The map m can be a partial function of the states. The states not specified are labeled with integers starting from i_0 in ascending order. If m is empty, then all states are relabeled with integers. Lastly, if *inf* is true then the tree $T_{\mathcal{A}}$ associated with the DFA is also relabeled, otherwise it is set as empty.

7.3. Operators

7.3.1. Hold

The DFA corresponding to a *hold* operator is constructed by [Algorithm 4](#). The algorithm takes as input an atomic proposition s in positive or negative form, a duration d , and the set of atomic propositions AP . The computed DFA has $d + 2$ states

Algorithm 2: $relabelTree(T, m, e)$.

Input: T – a tree structure
Input: m – (complete) relabeling mapping
Input: e – boolean, true if m maps states to sets of states

```

1 if  $\neg e$  then  $m(s) \leftarrow \{m(s)\}, \forall s$ 
2  $stack \leftarrow [T]$ 
3 while  $stack \neq []$  do
4    $t \leftarrow stack.pop()$ 
5    $t.I \leftarrow \bigcup_{s \in t.I} m(s)$ 
6    $t.F \leftarrow \bigcup_{s \in t.F} m(s)$ 
7   if  $op = \vee$  then
8      $B, L, R \leftarrow t.choices$ 
9      $B' \leftarrow \bigcup_{(s_B, \sigma) \in B} \{(s, \sigma) \mid s \in m(s_B)\}$ 
10     $L' \leftarrow \bigcup_{(s_L, \sigma) \in L} \{(s, \sigma) \mid s \in m(s_L)\}$ 
11     $R' \leftarrow \bigcup_{(s_R, \sigma) \in R} \{(s, \sigma) \mid s \in m(s_R)\}$ 
12     $t.choices \leftarrow (B', L', R')$ 
13 if  $t.left \neq \emptyset$  then  $stack.push(t.left)$ 
14 if  $t.right \neq \emptyset$  then  $stack.push(t.right)$ 

```

Algorithm 3: $relabel(\mathcal{A}, m, i_0)$.

Input: $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta, F_{\mathcal{A}})$ – a DFA
Input: m – (partial) relabeling mapping
Input: i_0 – start labeling index
Output: the relabeled DFA

```

1 for  $s \in S_{\mathcal{A}}$  s.t.  $\nexists m(s)$  do
2    $m(s) \leftarrow i_0$ 
3    $i_0 \leftarrow i_0 + 1$ 
4  $S'_{\mathcal{A}} \leftarrow \{m(s) \mid s \in S_{\mathcal{A}}\}$ 
5  $\delta' \leftarrow \{m(s) \xrightarrow{\sigma}_{\mathcal{A}} m(s') \mid s \xrightarrow{\sigma}_{\mathcal{A}} s'\}$ 
6  $F'_{\mathcal{A}} \leftarrow \{m(s) \mid s \in F_{\mathcal{A}}\}$ 
7 if  $inf$  then  $T'_{\mathcal{A}} \leftarrow relabelTree(T_{\mathcal{A}}, m)$ 
8 else  $T'_{\mathcal{A}} \leftarrow \emptyset$ 
9 return  $(S'_{\mathcal{A}}, m(s_0), \Sigma, \delta', F'_{\mathcal{A}}, T'_{\mathcal{A}})$ 

```

(line 1) that are connected in series as follows: (i) if s is in positive form then the states are connected by all transitions activated by symbols which contain s (lines 2–4); and (ii) if s is in negative form then the states are connected by all transitions activated by symbols which do not contain s (lines 5–7). Lastly, if inf is true, a new leaf node is created (line 8).

Algorithm 4: $Q_H(p, d, AP)$.

Input: $p \in \{s, \neg s\}$, $s \in AP$
Input: d – hold duration
Input: AP – set of atomic propositions
Output: DFA corresponding to $H^d p$

```

1  $S \leftarrow \{0, \dots, d+1\}$ 
2 if  $p = s$  then
3    $\Sigma_s \leftarrow 2^{AP} \setminus 2^{(AP \setminus \{s\})}$ 
4    $\delta \leftarrow \{i \xrightarrow{\sigma}_{\mathcal{A}} (i+1) \mid i \in \{0, \dots, d\}, \sigma \in \Sigma_s\}$ 
5 else
6    $\Sigma_{\neg s} \leftarrow 2^{(AP \setminus \{s\})}$ 
7    $\delta \leftarrow \{i \xrightarrow{\sigma}_{\mathcal{A}} (i+1) \mid i \in \{0, \dots, d\}, \sigma \in \Sigma_{\neg s}\}$ 
8 if  $inf$  then  $T \leftarrow tree(H^d, \emptyset, \emptyset, \{0\}, \{d+1\})$ 
9 else  $T \leftarrow \emptyset$ 
10 return  $(S, 0, 2^{AP}, \delta, \{d+1\}, T)$ 

```

7.3.2. Conjunction and disjunction

The construction for conjunction and disjunction operations is based on the synchronous product construction and is similar to the standard one [17]. However, Q_{\wedge} and Q_{\vee} produce strict DFAs, which only have one accepting state. Both algorithms recursively construct the product automaton starting from the initial composite state. In the following, we describe the details of the algorithms separately.

Conjunction: The DFA corresponding to the conjunction operation is constructed by Algorithm 5. The procedure is recursive and the synchronization condition, i.e., the transition relation, is the following: given two composite states (s_1, s_2)

and (s'_1, s'_2) , there exists a transition from the first state to the second state if there exists a symbol σ such that: (i) there exists pairwise transitions enabled by σ in the two automata (lines 9–11), i.e., $s_1 \xrightarrow{\sigma}_{\mathcal{A}_1} s'_1$ and $s_2 \xrightarrow{\sigma}_{\mathcal{A}_2} s'_2$; (ii) one automaton reached a final state and the other has a transition enabled by σ (lines 5–8), i.e., either (a) $s_1 = s'_1 = s_{f_1}$ and $s_2 \xrightarrow{\sigma}_{\mathcal{A}_2} s'_2$, or (b) $s_1 \xrightarrow{\sigma}_{\mathcal{A}_1} s'_1$ and $s_2 = s'_2 = s_{f_2}$. The first case covers the synchronous execution (simulation) of both \mathcal{A}_1 and \mathcal{A}_2 when a symbol is encountered. The second case corresponds to the situation when the two automata require words of different sizes to accept an input. A simple example of this case is the DFA encoding $H^2A \wedge H^3B$ and the input word $\{A, B\}, \{A, B\}, \{A, B\}, \{B\}$, which clearly satisfies the TWTL formula.

Algorithm 5: $\varrho_{\wedge}(\mathcal{A}_1, \mathcal{A}_2)$.

Input: $\mathcal{A}_1 = (S_{\mathcal{A}_1}, s_{01}, \Sigma, \delta_1, \{s_{f_1}\}, T_{\mathcal{A}_1})$ – left DFA
Input: $\mathcal{A}_2 = (S_{\mathcal{A}_2}, s_{02}, \Sigma, \delta_2, \{s_{f_2}\}, T_{\mathcal{A}_2})$ – right DFA
Output: DFA corresponding to $\mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$

```

1  $S \leftarrow \{(s_{01}, s_{02})\}, E \leftarrow \emptyset$ 
2  $stack \leftarrow [(s_{01}, s_{02})]$ 
3 while  $stack \neq []$  do
4    $s = (s_1, s_2) \leftarrow stack.pop()$ 
5   if  $s_1 = s_{f_1}$  then
6      $S_n \leftarrow \{(s_1, s'_2), \sigma \mid s_2 \xrightarrow{\sigma}_{\mathcal{A}_2} s'_2\}$ 
7   else if  $s_2 = s_{f_2}$  then
8      $S_n \leftarrow \{(s'_1, s_2), \sigma \mid s_1 \xrightarrow{\sigma}_{\mathcal{A}_1} s'_1\}$ 
9   else
10     $S_n \leftarrow \{(s'_1, s'_2), \sigma \mid \exists \sigma \in \Sigma \text{ s.t.}$ 
11       $(s_1 \xrightarrow{\sigma}_{\mathcal{A}_1} s'_1) \wedge (s_2 \xrightarrow{\sigma}_{\mathcal{A}_2} s'_2)\}$ 
12     $E \leftarrow E \cup \{(s, \sigma, s') \mid (s', \sigma) \in S_n\}$ 
13     $S' \leftarrow \{s' \mid \exists \sigma \in \Sigma \text{ s.t. } (s', \sigma) \in S_n\}$ 
14     $stack.extend(S' \setminus S)$ 
15     $S \leftarrow S \cup S'$ 
16  $m_L = \{(u, \{(u, v) \in S_{\mathcal{A}_1}) \mid u \in S_{\mathcal{A}_1}\}$ 
17  $m_R = \{(v, \{(u, v) \in S_{\mathcal{A}_1}) \mid v \in S_{\mathcal{A}_2}\}$ 
18  $T_{\mathcal{A}} \leftarrow tree(\wedge, relabelTree(T_{\mathcal{A}_1}, m_L, \top),$ 
19    $relabelTree(T_{\mathcal{A}_2}, m_R, \top), \{(s_{01}, s_{02})\}, \{(s_{f_1}, s_{f_2})\})$ 
20  $\mathcal{A} \leftarrow (S, (s_{01}, s_{02}), \Sigma, E, \{(s_{f_1}, s_{f_2})\}, T_{\mathcal{A}})$ 
21 return  $relabel(\mathcal{A}, \emptyset, 0)$ 

```

Note that [Algorithm 5](#) generates only composite states which are reachable from the initial composite state (s_{01}, s_{02}) . The resulting automaton has a single final state (s_{f_1}, s_{f_2}) which captures the fact that both automata must accept the input word in order for the product automaton to accept it.

After the automaton is constructed, the corresponding tree is created (lines 16–19). The child subtrees are taken from \mathcal{A}_1 and \mathcal{A}_2 , and relabeled. The relabeling mapping expands each state s to the set of all composite states, which have s as the first or second component corresponding to whether s is a state of the left or right automaton, respectively.

Disjunction: The disjunction operation is translated using [Algorithm 6](#). The first step of the algorithm is to add a trap state in each of the two automata \mathcal{A}_1 and \mathcal{A}_2 (line 1). All states of an automaton, except the final state, are connected via blocking symbols to the trap state \bowtie (lines 3–4). The trap state has self-transitions for all symbols. Afterwards, the algorithm creates the synchronous product automaton in the same way as for the conjunction operation (lines 4–13). However, in this case, we do not need to treat composite states that contain a final state of one of the two automata separately. This follows from the semantics of the disjunction operation, which accepts a word as soon as at least one automaton accepts the word.

In the standard construction [\[17\]](#), the resulting automaton would have multiple final states, which are computed in line 17. However, because final states do not have outgoing transitions, we can merge all final states and obtain an automaton with only one final state (lines 17–20). The composite trap state is also removed from the set of states (line 18).

The annotation tree is created similarly to the conjunction case (lines 21–24). However, for the disjunction case, we add additional information on the automaton. This information $T.choices$ is used in latter algorithm to determine if a word has satisfied the left, right, or both sub-formulae corresponding to the disjunction formula. This is done by partitioning the transitions incoming into final states (line 14–16) and storing this partition in the associated tree node (line 25). Note that only the start state and the symbol of each transition is stored in the partition sets and these are well defined, because the DFAs are deterministic.

7.3.3. Concatenation

The algorithm to compute an automaton accepting the concatenation language of two languages is shown in [Algorithm 7](#). The special structure of the unambiguous languages, see [Sec. 6](#) for details, admits a particularly simple and intuitive construction procedure. The composite automaton is obtained by identifying the final state of left automaton \mathcal{A}_1 with the initial state of the right automaton \mathcal{A}_2 .

Algorithm 6: $Q_{\vee}(\mathcal{A}_1, \mathcal{A}_2)$.

Input: $\mathcal{A}_1 = (S_{\mathcal{A}_1}, s_{01}, \Sigma, \delta_1, \{s_{f1}\}, T_{\mathcal{A}_1})$ – left DFA
Input: $\mathcal{A}_2 = (S_{\mathcal{A}_2}, s_{02}, \Sigma, \delta_2, \{s_{f2}\}, T_{\mathcal{A}_2})$ – right DFA
Output: DFA corresponding to $\mathcal{L}(\mathcal{A}_1) \cup \mathcal{L}(\mathcal{A}_2)$

- 1 $S'_{\mathcal{A}_1} \leftarrow S_{\mathcal{A}_1} \cup \{\perp\}$, $S'_{\mathcal{A}_2} \leftarrow S_{\mathcal{A}_2} \cup \{\perp\}$
- 2 $\delta'_1 \leftarrow \delta_1 \cup \{(s, \sigma, \perp) \mid s \in S'_{\mathcal{A}_1} \setminus \{s_{f1}\}, \sigma \in \Sigma, \nexists \delta_1(s, \sigma)\}$
- 3 $\delta'_2 \leftarrow \delta_2 \cup \{(s, \sigma, \perp) \mid s \in S'_{\mathcal{A}_2} \setminus \{s_{f2}\}, \sigma \in \Sigma, \nexists \delta_2(s, \sigma)\}$
- 4 $S \leftarrow \{(s_{01}, s_{02})\}$, $E \leftarrow \emptyset$
- 5 $stack \leftarrow [(s_{01}, s_{02})]$
- 6 **while** $stack \neq []$ **do**
- 7 $s = (s_1, s_2) \leftarrow stack.pop()$
- 8 $S_n \leftarrow \{(s'_1, s'_2), \sigma \mid \exists \sigma \in \Sigma \text{ s.t.}$
- 9 $(s'_1 = \delta'_1(s_1, \sigma)) \wedge (s'_2 = \delta'_2(s_2, \sigma))\}$
- 10 $E \leftarrow E \cup \{(s, \sigma, s') \mid (s', \sigma) \in S_n\}$
- 11 $S' \leftarrow \{s' \mid \exists \sigma \in \Sigma \text{ s.t. } (s', \sigma) \in S_n\}$
- 12 $stack.extend(S' \setminus S)$
- 13 $S \leftarrow S \cup S'$
- 14 $B \leftarrow \{(s, \sigma) \mid \exists \sigma \text{ s.t. } (s, \sigma, (s_{f1}, s_{f2})) \in E\}$
- 15 $L \leftarrow \{(s, \sigma) \mid \exists s_2 \neq s_{f2}, \exists \sigma \text{ s.t. } (s, \sigma, (s_{f1}, s_2)) \in E\}$
- 16 $R \leftarrow \{(s, \sigma) \mid \exists s_1 \neq s_{f1}, \exists \sigma \text{ s.t. } (s, \sigma, (s_1, s_{f2})) \in E\}$
- 17 $F \leftarrow \{(s_1, s_2) \in S \mid (s_1 = s_{f1}) \vee (s_2 = s_{f2})\}$
- 18 $S \leftarrow S \setminus (F \cup \{(s, \sigma, \perp)\})$
- 19 $E \leftarrow E \setminus \{(s, \sigma, s') \in E \mid s' \in F\}$
- 20 $E \leftarrow E \cup \{(s, \sigma, (s_{f1}, s_{f2})) \mid (s, \sigma) \in B \cup L \cup R\}$
- 21 $m_L = \{(u, \{(u, v) \in S_{\mathcal{A}_1}\}) \mid u \in S_{\mathcal{A}_1}\}$
- 22 $m_R = \{(v, \{(u, v) \in S_{\mathcal{A}_1}\}) \mid v \in S_{\mathcal{A}_2}\}$
- 23 $T_{\mathcal{A}} \leftarrow tree(\vee, relabelTree(T_{\mathcal{A}_1}, m_L, \top),$
- 24 $relabelTree(T_{\mathcal{A}_2}, m_R, \top), \{(s_{01}, s_{02})\}, \{(s_{f1}, s_{f2})\})$
- 25 $T_{\mathcal{A}}.choices \leftarrow (B, L, R)$
- 26 $\mathcal{A} \leftarrow (S, (s_{01}, s_{02}), \Sigma, E, \{(s_{f1}, s_{f2})\}, T_{\mathcal{A}})$
- 27 **return** $relabel(\mathcal{A}, \emptyset, 0)$

Algorithm 7: $Q_{\cdot}(\mathcal{A}_1, \mathcal{A}_2)$.

Input: $\mathcal{A}_1 = (S_{\mathcal{A}_1}, s_{01}, \Sigma, \delta_1, \{s_{f1}\}, T_{\mathcal{A}_1})$ – left DFA
Input: $\mathcal{A}_2 = (S_{\mathcal{A}_2}, s_{02}, \Sigma, \delta_2, \{s_{f2}\}, T_{\mathcal{A}_2})$ – right DFA
Output: DFA corresponding to $\mathcal{L}(\mathcal{A}_1) \cdot \mathcal{L}(\mathcal{A}_2)$

- 1 $\mathcal{A}_1 \leftarrow relabel(\mathcal{A}_1, \emptyset, 0)$
- 2 $\mathcal{A}_2 \leftarrow relabel(\mathcal{A}_2, \{(s_{02}, s_{f1})\}, |S_{\mathcal{A}_1}|)$
- 3 **if** inf **then** $T \leftarrow tree(\cdot, T_{\mathcal{A}_1}, T_{\mathcal{A}_2}, \{s_{01}\}, \{s_{f2}\})$
- 4 **else** $T \leftarrow \emptyset$
- 5 **return** $(S_{\mathcal{A}_1} \cup S_{\mathcal{A}_2}, s_{01}, \Sigma, \delta_1 \cup \delta_2, \{s_{f2}\}, T)$

7.3.4. Within

There are two algorithms used to construct a DFA associated with a *within* operator, [Algorithm 8](#) and [Algorithm 9](#) correspond to the relaxed and normal construction (lines 6–9 of [Algorithm 1](#)).

Relaxed within: The construction procedure [Algorithm 8](#) is as follows: starting from the DFA corresponding to the enclosed formula, all states are connected via blocking symbols to the initial state (lines 3–4). The last step is to create a number of a states connected in sequence for all symbols, similarly to [Algorithm 4](#), and connecting the a -th state to the initial state also for all symbols (lines 5–8).

Connecting all states to the initial state represents a restart of the automaton in case a blocking symbol was encountered. Thus, the resulting automaton offers infinite retries for a word to satisfy the enclosed formula. The a states added before the initial state represent a delay of length a for the start of the tracking of the satisfaction of the enclosed formula. Note that the procedure and resulting automaton do not depend on the upper bound b .

Normal within: The algorithm for the normal case builds upon [Algorithm 8](#). In this case the construction procedure [Algorithm 9](#) must take into account the upper time bound b . Similarly to the relaxed case, we need to restart the automaton of the when a blocking symbol is encountered. However, there are two major differences: (i) the automaton must track the number of restarts, because there are only a finite number of tries depending on the deadline b , and (ii) the automaton \mathcal{A} may need to be truncated for the last restart retries, i.e., all paths must have a length of at most a given length, in order to ensure that the satisfaction is realized before the upper time limit b .

In [Algorithm 9](#), first the maximum number of restarts p is computed in lines 1–2. Then, p DFAs are created (lines 3–12), which correspond to the relabeled and truncated copies of \mathcal{A} , see [Algorithm 10](#), and their union is computed iteratively. The truncation bound is computed as the remaining time units until the limit b is reached. The final state is always labeled with -1 (line 7) and, therefore, the resulting DFA has exactly one final state. Next, the restart transitions are added (lines 13–18).

Algorithm 8: $\varrho_{\infty}(\mathcal{A}, a, b)$.

Input: $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta, \{s_f\}, T_{\mathcal{A}})$ – child DFA
Input: a – lower bound of time-window
Input: b – upper bound of time-window
Output: computed DFA

```

1  $\mathcal{A} \leftarrow \text{relabel}(\mathcal{A}, \emptyset, 0)$ 
2  $S \leftarrow \emptyset, E \leftarrow \emptyset$ 
3 for  $s \in S_{\mathcal{A}} \setminus \{s_f\}$  do
4    $E \leftarrow E \cup \{(s, \sigma, s_0) \mid \nexists s' = \delta(s, \sigma)\}$ 
5 if  $a > 0$  then
6    $S \leftarrow \{|S_{\mathcal{A}}|, \dots, |S_{\mathcal{A}}| + a - 1\}$ 
7    $E \leftarrow E \cup \{(i, \sigma, i+1) \mid i \in S \setminus \{|S_{\mathcal{A}}| + a - 1\}, \sigma \in \Sigma\}$ 
8    $E \leftarrow E \cup \{(|S_{\mathcal{A}}| + a - 1, \sigma, s_0) \mid \sigma \in \Sigma\}$ 
9  $T \leftarrow \text{tree}([\infty^{a,b}], T_{\mathcal{A}}, \emptyset, \{|S_{\mathcal{A}}|\}, \{s_f\})$ 
10 return  $(S_{\mathcal{A}} \cup S, |S_{\mathcal{A}}|, \Sigma, \delta \cup E, \{s_f\}, T)$ 

```

Algorithm 9: $\varrho_{[\]}(\mathcal{A}, a, b)$.

Input: $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta, \{s_f\}, T_{\mathcal{A}})$ – child DFA
Input: a – lower bound of time-window
Input: b – upper bound of time-window
Output: computed DFA

```

1  $l \leftarrow \text{Dijkstra}(\mathcal{A}, s_0, s_f)$ 
2  $p \leftarrow b - a - l + 2$ 
3  $l \leftarrow []$  // list
4  $n \leftarrow 0$ 
5  $\mathcal{A}_r \leftarrow (S_{\mathcal{A}_r} = \emptyset, \infty, \Sigma, \delta_r = \emptyset, \emptyset, \emptyset)$ 
6 for  $k \in \{1, \dots, p\}$  do // mark final state
7    $m \leftarrow \{s_f, -1\}$ 
8    $\mathcal{A}_a \leftarrow \text{relabel}(\mathcal{A}, m, n)$ 
9    $\mathcal{A}_t \leftarrow \text{truncate}(\mathcal{A}_a, b - a + 2 - k)$ 
10   $\mathcal{A}_r \leftarrow (S_{\mathcal{A}_r} \cup S_{\mathcal{A}_t}, \infty, \Sigma, \delta_r \cup \delta_t, \{-1\}, \emptyset)$ 
11   $l \leftarrow l + |S_{\mathcal{A}_t}|$ 
12   $n \leftarrow n + |S_{\mathcal{A}_t}|$ 
13  $S_c \leftarrow \{l[0]\}, E \leftarrow \emptyset$ 
14 for  $s_r \in l[1:]$  do
15   $S_n \leftarrow \emptyset$ 
16  for  $s \in S_c \setminus \{-1\}$  do
17     $E \leftarrow E \cup \{(s, \sigma, s_r) \mid \sigma \in \Sigma \text{ s.t. } \nexists \delta_r(s, \sigma)\}$ 
18   $S_c \leftarrow S_c \cup \{s_r\}$ 
19  $S \leftarrow \emptyset$ 
20 if  $a > 0$  then
21   $S \leftarrow \{|S_{\mathcal{A}_r}|, \dots, |S_{\mathcal{A}_r}| + a - 1\}$ 
22   $E \leftarrow E \cup \{(i, \sigma, i+1) \mid i \in S \setminus \{|S_{\mathcal{A}_r}| + a - 1\}, \sigma \in \Sigma\}$ 
23   $E \leftarrow E \cup \{(|S_{\mathcal{A}_r}| + a - 1, \sigma, s_0) \mid \sigma \in \Sigma\}$ 
24 return  $(S_{\mathcal{A}_r} \cup S, l[0], \Sigma, \delta_r \cup E, \{-1\}, \emptyset)$ 

```

Note that the transitions, enabled by blocking symbols, lead to initial states of the proper restart automaton. For example, if a blocking symbol was encountered after two symbols, then the restart transition (if it exists) leads to the initial state of the fourth copy of the automaton. Lastly, a delay of a time units is added before the initial state of the automaton similar to the relaxed case.

7.3.5. *Truncate*

Algorithm 10 takes as input a DFA \mathcal{A} and a cutoff bound l and returns a version of \mathcal{A} with all paths guaranteed to have length at most l . The algorithm is based on a breath-first search and returns a strict DFA.

7.4. *Correctness*

The following theorems show that the proposed algorithms for translating TWTL formulae to (normal or annotated) automata are correct.

Theorem 7.1. *If ϕ is a TWTL formula satisfying [Assumption 1](#) and the global parameter `inf` is true, then [Algorithm 1](#) generates a DFA \mathcal{A}_{∞} such that $\mathcal{L}(\mathcal{A}_{\infty}) = \mathcal{L}(\phi(\infty))$.*

Algorithm 10: $\text{truncate}(\mathcal{A}, l)$.

Input: $\mathcal{A} = (S_{\mathcal{A}}, s_0, \Sigma, \delta, \{s_f\}, T_{\mathcal{A}})$ – a DFA
Input: l – cutoff value
Output: computed DFA

```

1  $S \leftarrow \{s_0\}$ 
2  $E \leftarrow \emptyset$ 
3  $L_n \leftarrow \{s_0\}$ 
4 for  $i \in \{1, \dots, l\}$  do
5    $L_c \leftarrow L_n$ 
6    $L_n \leftarrow \emptyset$ 
7   for  $s \in L_c$  do
8     for  $(s_c, \sigma_c) \in \{(s', \sigma) \mid \exists \sigma \in \Sigma \text{ s.t. } s \xrightarrow{\sigma} s'\}$  do
9        $E \leftarrow E \cup (s, \sigma_c, s_c)$ 
10      if  $s_c \notin S$  then
11         $S \leftarrow S \cup \{s_c\}$ 
12         $L_n \leftarrow L_n \cup \{s_c\}$ 
13  $\mathcal{A}_t = (S_{\mathcal{A}}, s_0, \Sigma, \delta \setminus E, \{s_f\}, T_{\mathcal{A}})$ 
14  $S_{\text{traps}} = \{s \in S_{\mathcal{A}} \mid \nexists \sigma \in \Sigma^* \text{ s.t. } s \xrightarrow{\sigma} s_f\}$ 
15 return  $(S_{\mathcal{A}} \setminus S_{\text{traps}}, s_0, \Sigma, \delta \setminus E, \{s_f\}, T_{\mathcal{A}})$ 

```

Proof. The proof follows by structural induction on $AST(\phi)$ and the properties of TWTL languages.

Before we proceed with the induction, notice that all construction algorithms associated with the operators of TWTL generate strict DFAs with only one final state without any outgoing transitions.

The base case corresponds to the leaf nodes of $AST(\phi)$ which are associated with *hold* operators, see Fig. 1, and follows by construction from Algorithm 4.

The induction hypothesis requires that the theorem holds for the DFAs returned by the recursion in Algorithm 1. In the case of the conjunction and disjunction operators, the property follows from the product construction method [17]. The theorem holds also for the concatenation operator, because: (a) the returned DFAs have one final state without any outgoing transitions, and (b) the languages corresponding to the two operand formulae are unambiguous. Thus, the correctness of the construction described in Algorithm 7 follows immediately from the unambiguity of the concatenation, see Definition 6.2. Lastly, the case of the *within* operator (relaxed form), follow from the Assumption 1. The *within* operator adds transitions to a DFA from each state to the initial state on all undefined symbols. In other words, the operator restarts the execution of a DFA from the initial state. If there are no disjunction operators, then going back to the initial state is the only correct choice. Otherwise, the information about which paths need to be restarted is lost, because the paths induced by a disjunction become indistinguishable from each other. Thus, it is not possible to properly add backward transitions to restart paths associated with the terms of the disjunction independently in the automaton as needed to capture the *within* operator. \square

Theorem 7.2. *If ϕ is a TWTL formula satisfying Assump. 1 and the global parameter inf is false, then Algorithm 1 generates DFA \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\phi)$.*

Proof. The proof is similar to that of Theorem 7.1 and is omitted for brevity. \square

7.5. Complexity

In this section, we review the complexity of the algorithms presented in the previous section for the construction of DFAs from TWTL formulae. The complexity of basic composition operations for incomplete and acyclic DFAs has been explored in [32,16,7,12,9]. Our construction algorithms differ from the ones in the literature because we specialized and optimized them to translate TWTL formulae and handle words over power sets of atomic propositions.

The complexity of the relabeling procedures are $O(|T|)$ and $O(|S_{\mathcal{A}}|)$ corresponding to Algorithm 2 and Algorithm 3, respectively. The complexity of the *hold* operator Algorithm 4 is $O(d \cdot 2^{|AP|})$. The construction algorithms for conjunction and disjunction Algorithm 5 and Algorithm 6 have the same complexity $O(|S_{\mathcal{A}_1}| \cdot |S_{\mathcal{A}_2}| \cdot 2^{|AP|})$, because these are based on the product automaton construction. Concatenation has complexity $O(|S_{\mathcal{A}_1}| + |S_{\mathcal{A}_2}|)$ due to the relabeling operations. Lastly, the *within* operation can be performed in $O(a \cdot 2^{|AP|} + |S_{\mathcal{A}}| \cdot 2^{|AP|})$ and $O(a \cdot 2^{|AP|} + b |S_{\mathcal{A}}| \cdot 2^{|AP|})$ for the infinity Algorithm 8 and the normal Algorithm 9 construction, respectively, where Algorithm 10 used in the normal construction procedure takes $O(|S_{\mathcal{A}}| \cdot 2^{|AP|})$. The overall translation algorithm Algorithm 1 takes at most $O(2^{|\phi|+|AP|})$.

It is very important to notice that the infinity construction does not depend on the deadline b , which makes the procedure more efficient than the normal construction.

7.6. Ambiguity checking

The following result establishes the complexity of checking the ambiguity of a language.

Proposition 7.3. *Checking if a language is unambiguous takes at most $O(2^{|\phi|+|AP|})$.*

Proof. First, note that the languages generated by *hold* operators are unambiguous, and the only operation that may create ambiguous languages is the disjunction operation. The claim follows from Proposition 6.3 and 6.5 for concatenation and conjunction, respectively. The case of the *within* operator follows from its definition, since a word satisfying the inner formula immediately implies the satisfaction of the overall formula.

To check for ambiguity we use Proposition 6.4 in Algorithm 6. At line 17, we check if the set of final states of the product automaton $F \subseteq \{(s_{f1}, \infty), (\infty, s_{f2}), (s_{f1}, s_{f2})\}$, which is equivalent to $\mathcal{L}(\phi_1) \cap P(\mathcal{L}(\phi_2)) = \emptyset$ and $\mathcal{L}(\phi_2) \cap P(\mathcal{L}(\phi_1)) = \emptyset$. Thus, the complexity of checking if the language of a TWTL formula is unambiguous is at most the same as computing the corresponding automaton. \square

8. Verification, synthesis, and learning algorithms

In this section, we will use the following notation. Let T be an annotation tree associated with a DFA. We denote by ϕ_T the TWTL formula corresponding to the tree T . Given a finite sequence $\mathbf{p} = p_0, \dots, p_n$, we denote the first and the last elements by $b(\mathbf{p}) = p_0$ and $e(\mathbf{p}) = p_n$, respectively.

Definition 8.1 (Primitive). Let ϕ be a TWTL formula. We say that ϕ is *primitive* if ϕ does not contain any *within* operators.

8.1. Compute temporal relaxation for a word

The automata construction presented in Sec. 7 can be used to compute the temporal relaxation of words with respect to TWTL formulae. Let ϕ be a TWTL formula and σ be a word. In this section, we show how to infer (synthesize) a set of temporal relaxations τ of the deadlines in ϕ such that σ satisfies $\phi(\tau)$ and $|\tau|_{TR}$ is minimized. Algorithm 11 computes the vector of temporal relaxations corresponding to each *within* operator. First, the annotated DFA \mathcal{A}_∞ is computed together with the associated annotation tree T (line 2). Next, additional annotations are added to the tree T using the *initTreeTR()* procedure (line 3). Each node corresponding to a *within* operation is assigned three variables $T.ongoing$, $T.done$ and $T.steps$, which track whether the processing of the operator is ongoing, done, and the number of steps to process the operator, respectively. The three variables are initialized to \perp , \perp , and -1 , respectively. Then, Algorithm 11 cycles through the symbols of the input word σ and updates the tree using *updateTree()* via Algorithm 12. Finally, the temporal relaxation vector is returned by the *evalTreeTR()* procedure via Algorithm 13.

Algorithm 11: *tr*(\cdot) – Compute temporal relaxation.

Input: σ a word over the alphabet 2^{AP}
Input: ϕ a TWTL formula
Output: τ^* – minimum maximal temporal relaxation
Output: τ – temporal relaxation vector

```

1 if  $\phi$  is primitive then return  $(-\infty, [])$ 
2  $\mathcal{A}_\infty, T \leftarrow \text{translate}(\phi; \text{inf} = \top)$ 
3 initTreeTR( $T$ )
4  $s_{prev} \leftarrow \perp; s_c \leftarrow s_0$ 
5 updateTreeTR( $T, s_c, s_{prev}, \emptyset, \emptyset$ )
6 for  $\sigma \in \sigma$  do
7   if  $s_c \in F_{\mathcal{A}_\infty}$  then break
8    $s_{prev} \leftarrow s_c$ 
9    $s_c \leftarrow \delta_{\mathcal{A}_\infty}(s_c, \sigma)$ 
10  updateTreeTR( $T, s_c, s_{prev}, \sigma, \emptyset$ )
11 return evalTreeTR( $T$ )

```

The tree is updated recursively in Algorithm 12. A *within* operator is marked as ongoing, i.e., $T.ongoing = \top$, when the current state is in the set of initial states associated with the operator (line 2). Similarly, when the current state is in the set of final states associated with the operator, the *within* operator is marked as done (lines 3–6), i.e. $T.done = \top$ and $T.ongoing = \perp$. The number of steps $T.steps$ of all ongoing *within* operators is incremented (line 7).

To enforce correct computation of the temporal relaxation with respect to the disjunction operators, Algorithm 12 keeps track of a set of constraints C . The set C is composed of state-symbol pairs, and is used to determine which of the two subformulae of a disjunction are satisfied by the input word (lines 12–17). To achieve this, we use the annotation variables $T.choices$ (see Algorithm 6), which capture both cases. For all other operators, the constraint sets are propagated unchanged (lines 8, 10, 11).

Finally, Algorithm 13 extracts the temporal relaxation from the annotation tree T after all symbols of the input word σ were processed. Algorithm 13 also computes the minimum maximum temporal relaxation value, which may be $-\infty$ if ϕ is primitive (line 1). The recursion in Algorithm 13 differs between disjunction and the other operators. One subformula is sufficient to hold to satisfy the formula associated with a disjunction operator. Thus, the optimal temporal relaxation is the minimum or maximum between the two optimal temporal relaxations of the subformulae for disjunction (line 12),

Algorithm 12: $updateTreeTR(\cdot)$.

```

Input:  $s_c$  – current state
Input:  $s_{prev}$  – previous state
Input:  $\sigma$  – current symbol in word
Input:  $C$  – set of constraints associated with the states

1 if  $T.op = [ ]^{[a,b]}$  then
2   if  $s_c \in T.I$  then  $T.ongoing \leftarrow \top$ 
3   if  $s_c \in T.F$  then
4     if  $(C = \emptyset) \vee (\sigma \subseteq C(s_{prev}))$  then
5        $T.ongoing \leftarrow \perp$ 
6        $T.done \leftarrow \top$ 
7   if  $T.ongoing$  then  $T.\tau \leftarrow T.\tau + 1$ 
8    $updateTreeTR(T.left, s_c, s_{prev}, \sigma, C)$ 
9 else
10  if  $T.op = \cdot$  then  $C_L \leftarrow \emptyset; C_R \leftarrow C$ 
11  else if  $T.op = \wedge$  then  $C_L \leftarrow C; C_R \leftarrow C$ 
12  else if  $T.op = \vee$  then
13     $C_L \leftarrow T.choices.L \cup T.choices.B$ 
14     $C_R \leftarrow T.choices.R \cup T.choices.B$ 
15    if  $C \neq \emptyset$  then
16       $C_L \leftarrow C \cap C_L$ 
17       $C_R \leftarrow C \cap C_R$ 
18   $updateTreeTR(T.left, s_c, s_{prev}, \sigma, C_L)$ 
19   $updateTreeTR(T.right, s_c, s_{prev}, \sigma, C_R)$ 

```

and conjunction and concatenation (line 13), respectively. Lines 15–16 of Algorithm 13 cover the cases involving primitive subformulae.

The complexity of Algorithm 11 is $O(2^{|\phi|+|AP|} + |\sigma| \cdot |\phi|)$, where the first term is the complexity of constructing \mathcal{A}_∞ in line 1 and the second term corresponds to the update of the tree for each symbol in σ and the final evaluation of the tree.

Algorithm 13: $evalTreeTR(\cdot)$.

```

Input:  $T$  – annotated tree
Output:  $\tau^*$  – minimum maximal temporal relaxation
Output:  $\tau$  – temporal relaxation vector

1 if  $\phi_T$  is primitive then return  $(-\infty, [ ])$ 
2 else if  $T.op = [\phi]^{[a,b]}$  then
3    $\tau_{ch}^*, \tau_{ch} = evalTreeTR(tree.left)$ 
4   if  $T.done = \top$  then
5     return  $(\max\{\tau_{ch}^*, T.steps - b\}, [\tau_{ch}, T.steps - b])$ 
6   else
7     return  $(-\infty, [\tau_{ch}, -\infty])$ 
8 else //  $\wedge, \vee$  or  $\cdot$ 
9    $\tau_L^*, \tau_L = evalTreeTR(tree.left)$ 
10   $\tau_R^*, \tau_R = evalTreeTR(tree.right)$ 
11  if  $(\tau_L^* \neq -\infty) \wedge (\tau_R^* \neq -\infty)$  then
12    if  $T.op = \vee$  then  $\tau^* \leftarrow \min\{\tau_L^*, \tau_R^*\}$ 
13    else  $\tau^* \leftarrow \max\{\tau_L^*, \tau_R^*\}$ 
14  else
15    if  $T.op = \wedge$  then  $\tau^* \leftarrow \max\{\tau_L^*, \tau_R^*\}$ 
16    else  $\tau^* \leftarrow -\infty$ 
17  return  $(\tau^*, [\tau_L, \tau_R])$ 

```

8.2. Control policy synthesis for a finite transition system

Let \mathcal{T} be a finite transition system, and ϕ a specification given as a TWTL formula. The procedure to synthesize an optimal control policy by minimizing the temporal relaxation has three steps:

1. constructing the annotated DFA \mathcal{A}_∞ corresponding to ϕ ,
2. constructing the synchronous product $\mathcal{P} = \mathcal{T} \times \mathcal{A}_\infty$ between the transition system \mathcal{T} and the annotated DFA \mathcal{A}_∞ ,
3. computing the optimal policy on \mathcal{P} using Algorithm 14 and generating the optimal trajectory of \mathcal{T} from the optimal trajectory of \mathcal{P} by projection,

where the synchronous product \mathcal{P} is defined as follows:

Definition 8.2 (Product automaton). Given a TS $\mathcal{T} = (X, x_0, \Delta, AP, h)$ and a DFA $\mathcal{A} = (S_{\mathcal{A}}, s_0, 2^{AP}, \delta_{\mathcal{A}}, F_{\mathcal{A}})$, their product automaton, denoted by $\mathcal{P} = \mathcal{T} \times \mathcal{A}$, is a tuple $\mathcal{P} = (S_{\mathcal{P}}, p_0, \Delta_{\mathcal{P}}, F_{\mathcal{P}})$ where:

- $p_0 = (x_0, s_0)$ is the initial state;
- $S_{\mathcal{P}} \subseteq X \times S_{\mathcal{A}}$ is a finite set of states that are reachable from the initial state: for every $(x^*, s^*) \in S_{\mathcal{P}}$, there exists a sequence of $\mathbf{x} = x_0 x_1 \dots x_n x^*$, with $x_k \rightarrow_{\mathcal{T}} x_{k+1}$ for all $0 \leq k < n$ and $x_n \rightarrow_{\mathcal{T}} x^*$, and a sequence $\mathbf{s} = s_0 s_1 \dots s_n s^*$ such that s_0 is the initial state of \mathcal{A} , $s_k \xrightarrow{h(x_{k+1})} s_{k+1}$ for all $0 \leq k < n$ and $s_n \xrightarrow{h(x^*)} s^*$;
- $\Delta_{\mathcal{P}} \subseteq S_{\mathcal{P}} \times S_{\mathcal{P}}$ is the set of transitions defined by: $((x, s), (x', s')) \in \Delta_{\mathcal{P}}$ iff $x \rightarrow_{\mathcal{T}} x'$ and $s \xrightarrow{h(x')} s'$;
- $F_{\mathcal{P}} = (X \times F_{\mathcal{A}}) \cap S_{\mathcal{P}}$ is the set of accepting states of \mathcal{P} .

A transition in \mathcal{P} is also denoted by $(x, s) \rightarrow_{\mathcal{P}} (x', s')$ if $((x, s), (x', s')) \in \Delta_{\mathcal{P}}$. A trajectory $\mathbf{p} = (x_0, s_0)(x_1, s_1) \dots$ of \mathcal{P} is an infinite sequence, where $(x_0, s_0) = p_0$ and $(x_k, s_k) \rightarrow_{\mathcal{P}} (x_{k+1}, s_{k+1})$ for all $k \geq 0$. A trajectory of $\mathcal{P} = \mathcal{T} \times \mathcal{A}$ is said to be accepting if and only if it ends in a state that belongs to the set of final states $F_{\mathcal{P}}$. It follows by construction that a trajectory $\mathbf{p} = (x_0, s_0)(x_1, s_1) \dots$ of \mathcal{P} is accepting if and only if the trajectory $s_0 s_1 \dots$ is accepting in \mathcal{A} . As a result, a trajectory of \mathcal{T} obtained from an accepting trajectory of \mathcal{P} satisfies the given specification encoded by \mathcal{A} . We denote the projection of a trajectory $\mathbf{p} = (x_0, s_0)(x_1, s_1) \dots$ onto \mathcal{T} by $\gamma_{\mathcal{T}}(\mathbf{p}) = x_0 x_1 \dots$.

Before we present the details of the proposed algorithm, we want to point out that completeness may be decided easily by using the product automaton \mathcal{P} . That is, testing if there exists a temporal relaxation such that a satisfying policy in \mathcal{T} may be synthesized can be performed very efficiently as shown by the following theorem.

Theorem 8.1. Let ϕ be a TWTL formula and \mathcal{T} be a finite transition system. Deciding if there exists a finite $\tau \in \mathbb{Z}^m$ and a trajectory \mathbf{x} of \mathcal{T} such that $\mathbf{o} \models \phi(\tau)$, can be performed in $O(|\Delta| \cdot |\delta_{\mathcal{A}_{\infty}}|)$, where m is the number of within operators in ϕ , \mathcal{A}_{∞} is the annotated DFA corresponding to ϕ , \mathbf{o} is the output trajectory induced by \mathbf{x} , and Δ and $\delta_{\mathcal{A}_{\infty}}$ are the sets of transitions of \mathcal{T} and \mathcal{A}_{∞} , respectively.

Remark 8.2. The complexity in Theorem 8.1 is independent of the deadlines of the within operators ϕ .

Proof. The result follows immediately from the construction of the product automaton \mathcal{P} by checking if the set of accepting states is empty. Because \mathcal{P} contains all reachable states, accepting states exists if and only if there is a trajectory \mathbf{x} of \mathcal{T} that generates an output word \mathbf{o} satisfying a relaxed specification $\phi(\tau)$, for some finite τ . If there are accepting states, then a trajectory \mathbf{x} can be computed using Dijkstra's algorithm on \mathcal{P} from the initial state to an accepting state, and then projecting onto \mathcal{T} . \square

Note that Dijkstra's algorithm may not necessarily provide an optimal trajectory of \mathcal{T} with respect to the minimum maximum temporal relaxation of the induced output word. Thus, we present a Dijkstra-based procedure to compute an optimal policy using the product automaton \mathcal{P} . The proposed solution is presented in Algorithm 14, which describes a recursive procedure over an annotated AST tree T .

The recursive procedure in Algorithm 14 has six cases. The first case (lines 1–3) corresponds to a primitive formula. In this case, there are no deadlines to relax since the formula does not contain any within operators. Thus, solutions (if any exist) can be computed using Dijkstra's algorithm. The next two cases treat the within operators. In the former case (lines 4–5), the enclosed formula is a primitive formula and the only deadline which must be optimized is the one associated with the current within operator. In the latter case (lines 7–10), the enclosed formula is not primitive. Therefore, there are multiple deadlines that must be considered. To optimize the temporal relaxation $|\cdot|_{TR}$, we take the maximum between the previous maximum temporal relaxation and the current temporal relaxation (line 10). The fourth case (lines 11–15) handles the concatenation operator. First, the paths and the corresponding temporal relaxations are computed for the left and the right subformulae in lines 12 and 13, respectively. Afterwards, the paths satisfying the left subformula are concatenated to the paths satisfying the right formula. However, the concatenation of paths p_L and p_R is restricted to pairs which have the following property: there exists a transition in \mathcal{P} between the last state of p_L and the first state in p_R . The temporal relaxation of the concatenation of two paths is the maximum between the temporal relaxations of the two paths (line 15). The next case is associated with the disjunction operator (lines 16–20). As in the concatenation case, first the paths satisfying the left M_L and the right M_R subformulae are computed in lines 17 and 18, respectively. The set corresponding to the disjunction of the two formulae is the union of the two sets because the paths must satisfy either one of the two subformulae. The temporal relaxation of a path p in the union is computed as follows (line 20): (a) if a path is only in the left, $\mathbf{p} \in M_L \setminus M_R$, or only in the right set, $\mathbf{p} \in M_R \setminus M_L$, then the temporal relaxation is $\tau_L^*[\mathbf{p}]$ or $\tau_R^*[\mathbf{p}]$, respectively; (b) the path is in both sets, $\mathbf{p} \in M_L \cap M_R$, then the temporal relaxations is the minimum of the two previously computed ones, $\min\{\tau_L^*[\mathbf{p}], \tau_R^*[\mathbf{p}]\}$. In the case (a), \mathbf{p} satisfies only one subformula and, therefore, only one temporal relaxation is available. In the case (b), \mathbf{p} satisfies both subformulae. Because only one is needed, the subformula that yields the minimum temporal relaxation is chosen, i.e., the minimum between the two temporal relaxations. The last case handles the conjunction operator (lines 21–25). As in the previous two cases, the paths satisfying the left and the right subformulae are computed first (lines 22–23). Then the intersection of the two sets is computed as the set of paths satisfying the conjunctions because the paths must satisfy both subformulae. The temporal relaxations of the paths in the intersections are computed as the maxima between the previously computed temporal relaxations for the left and the right subformulae.

Algorithm 14: Policy synthesis – $\text{policy}(T, \mathcal{P})$.

Input: T – the annotation AST tree
Input: \mathcal{P} – product automaton

- 1 **if** ϕ_T is primitive **then**
- 2 $M = \{\mathbf{p} \mid b(\mathbf{p}) \in T.I, e(\mathbf{p}) \in T.F\}$
- 3 $\tau^*(\mathbf{p}) = -\infty, \forall \mathbf{p} \in M$
- 4 **else if** $T.op = []^{[a,b]} \wedge \phi_{T.left}$ is primitive **then**
- 5 $M = \{\mathbf{p} \mid b(\mathbf{p}) \in T.I, e(\mathbf{p}) \in T.F\}$
- 6 $\tau^*(\mathbf{p}) = |\mathbf{p}| - b, \forall \mathbf{p} \in M$
- 7 **else if** $T.op = []^{[a,b]} \wedge \phi_{T.left}$ is not primitive **then**
- 8 $M_{ch}, \tau_{ch}^{max} = \text{policy}(T.left, \mathcal{P})$
- 9 $M = \{p_i \xrightarrow{a} p \xrightarrow{*} p' \mid p_i \in T.I, p \xrightarrow{*} p' \in M_{ch}\}$
- 10 $\tau^*(\mathbf{p}) = \max\{|\mathbf{p}| - b, \tau_{ch}^*(\mathbf{p})\}, \forall \mathbf{p} \in M$
- 11 **else if** $T.op = \cdot$ **then**
- 12 $M_L, \tau_L^* = \text{policy}(T.left, \mathcal{P})$
- 13 $M_R, \tau_R^* = \text{policy}(T.right, \mathcal{P})$
- 14 $M = \{\mathbf{p}_1 \cdot \mathbf{p}_2 \mid \mathbf{p}_1 \in M_L, \mathbf{p}_2 \in M_R, e(\mathbf{p}_1) \rightarrow_{\mathcal{P}} b(\mathbf{p}_2)\}$
- 15 $\tau^*(\mathbf{p}) = \max\{\tau_L^*(\mathbf{p}), \tau_R^*(\mathbf{p})\}, \forall \mathbf{p} \in M$
- 16 **else if** $T.op = \vee$ **then**
- 17 $M_L, \tau_L^* = \text{policy}(T.left, \mathcal{P})$
- 18 $M_R, \tau_R^* = \text{policy}(T.right, \mathcal{P})$
- 19 $M = M_L \cup M_R$
- 20 $\tau^*(\mathbf{p}) = \begin{cases} \tau_L^*(\mathbf{p}) & \mathbf{p} \in M \setminus M_R \\ \tau_R^*(\mathbf{p}) & \mathbf{p} \in M \setminus M_L \\ \min\{\tau_L^*(\mathbf{p}), \tau_R^*(\mathbf{p})\} & \mathbf{p} \in M_L \cap M_R \end{cases}$
- 21 **else if** $T.op = \wedge$ **then**
- 22 $M_L, \tau_L^* = \text{policy}(T.left, \mathcal{P})$
- 23 $M_R, \tau_R^* = \text{policy}(T.right, \mathcal{P})$
- 24 $M = M_L \cap M_R$
- 25 $\tau^*(\mathbf{p}) = \max\{\tau_L^*(\mathbf{p}), \tau_R^*(\mathbf{p})\}, \forall \mathbf{p} \in M$
- 26 **return** (M, τ^*)

Note that considering primitive formulae in Algorithm 14, instead of traversing the AST all the way to the leaves, optimizes the running time and the level of recursion of the algorithm.

A very important property of Algorithm 14 is that its complexity does not depend on the deadlines associated with the *within* operators of the TWTL specification formula ϕ . This is an immediate consequence of the DFA construction proposed in Sec 7. Moreover, it follows from Remark 4.3 that the completeness with respect to ϕ (unrelaxed) may also be decided independently of the values of the deadline values. Formally, we have the following results.

Theorem 8.3. *Let ϕ be a TWTL formula and \mathcal{T} be a finite transition system. Synthesizing a trajectory \mathbf{x} of \mathcal{T} such that $\mathbf{o} \models \phi(\tau)$ and $|\tau|_{TR}$ is minimized can be performed in $O(|\phi| \cdot |\Delta| \cdot |\delta_{\mathcal{A}_\infty}|)$, where $\tau \in \mathbb{Z}^m$, m is the number of within operators in ϕ , \mathcal{A}_∞ is the annotated DFA corresponding to $\phi(\infty)$, \mathbf{o} is the output trajectory induced by \mathbf{x} , and Δ and $\delta_{\mathcal{A}_\infty}$ are the sets of transitions of \mathcal{T} and \mathcal{A}_∞ , respectively.*

Proof. The worst-case complexity of Algorithm 14 is achieved when the TWTL formula ϕ has the form of primitive formulae enclosed by *within* operators and then composed by either the conjunction, disjunction, and concatenation operators.

The recursive algorithm stops when it encounters the primitive formulae and executes Dijkstra's algorithm that takes at most $O(|\Delta \mathcal{P}|) = O(|\Delta| \cdot |\delta_{\mathcal{A}_\infty}|)$ time. Since the recursion is performed with respect to an AST T of ϕ , the algorithm processes each operator only once. The complexity bound follows because the size of the set of paths M returned by the algorithm is at most the sum of the sized of the sets corresponding to the left and the right sets M_L and M_R , respectively. Thus, we obtain the bound $O(|\phi| \cdot |\Delta| \cdot |\delta_{\mathcal{A}_\infty}|)$ by summing up the time complexity over all nodes of T . \square

Corollary 8.4. *Let ϕ be a TWTL formula and \mathcal{T} be a finite transition system. Deciding if there exists a trajectory \mathbf{x} of \mathcal{T} such that $\mathbf{o} \models \phi$ can be performed in $O(|\phi| \cdot |\Delta| \cdot |\delta_{\mathcal{A}_\infty}|)$, where \mathcal{A}_∞ is the annotated DFA corresponding to ϕ , \mathbf{o} is the output trajectory induced by \mathbf{x} , and Δ and $\delta_{\mathcal{A}_\infty}$ are the sets of transitions of \mathcal{T} and \mathcal{A}_∞ , respectively.*

Proof. It follows from Theorem 8.3 and Remark 4.3. \square

8.3. Verification

The procedure described in Algorithm 15 solves the verification problem of a transition system \mathcal{T} against all relaxed versions of a TWTL specification First, the annotated DFA \mathcal{A}_∞ corresponding to ϕ is computed (line 1). Then a trap state \bowtie

is added in line 2 (see Algorithm 6 for details). Note that the final state is not connected to trap state. The transition system \mathcal{T} is composed with the DFA \mathcal{A}_∞ to produce the product automaton \mathcal{P} (line 3).

Lastly, it is checked if all trajectories of \mathcal{P} reach the final state in finite time (line 4), i.e., satisfy a relaxation of ϕ . The condition in line 4 ensures that: (i) there are final states; (ii) all paths are finite, i.e., \mathcal{P} is a DAG; and (iii) the only allowed sink states are the final states, i.e., the out degree $\deg(v)$ of all non-final states v of \mathcal{P} is positive.

Algorithm 15: Verification.

Input: \mathcal{T} – transition system
Input: ϕ – TWTL specification
Output: Boolean value

```

1  $\mathcal{A}_\infty \leftarrow \text{translate}(\phi; \text{inf} = \top)$ 
2 add trap state  $\triangleright$  to  $\mathcal{A}_\infty$ 
3  $\mathcal{P} \leftarrow \mathcal{T} \times \mathcal{A}_\infty$ 
4 return  $F_{\mathcal{P}} \neq \emptyset \wedge \text{isDAG}(\mathcal{P}) \wedge \left( \bigwedge_{p \in S_{\mathcal{P}} \setminus F_{\mathcal{P}}} \deg(v) > 0 \right)$ 
```

8.4. Learning deadlines from data

In this section, we present a simple heuristic procedure to infer deadlines from a finite set of labeled traces such that the misclassification rate is minimized. Let ϕ be a TWTL formula and \mathcal{L}_p and \mathcal{L}_n be two finite sets of words labeled as positive and negative examples, respectively. The misclassification rate is $|\{w \in \mathcal{L}_p \mid w \not\models \phi(\tau)\}| + |\{w \in \mathcal{L}_n \mid w \models \phi(\tau)\}|$, where $\phi(\tau)$ is a feasible τ -relaxation of ϕ . The terms of the misclassification rate are the false negative and false positive rates, respectively.

The procedure presented in Algorithm 16 uses Algorithm 11 to compute the tightest deadlines for each trace. Then each deadline is determined in a greedy way such that the misclassification rate is minimized. The heuristic in Algorithm 11 is due to the fact that each deadline is considered separately from the others. However, the deadlines are not independent with respect to the minimization of the misclassification rate.

Notice that the algorithm constructs \mathcal{A}_∞ only once at line 1. Then the automaton is used in the $tr(\cdot)$ function to compute the temporal relaxation of each trace, lines 2–3. Thus, the procedure avoids building \mathcal{A}_∞ for each trace.

Algorithm 16: Parameter learning.

Input: \mathcal{L}_p – set of positive traces
Input: \mathcal{L}_n – set of negative traces
Input: ϕ – template TWTL formula
Output: d – the vector of deadlines

```

1  $\mathcal{A}_\infty \leftarrow \text{translate}(\phi; \text{inf} = \top)$ 
2  $D_p \leftarrow \{tr(p, \mathcal{A}_\infty) + \mathbf{b} \mid p \in \mathcal{L}_p\}$ 
3  $D_n \leftarrow \{tr(p, \mathcal{A}_\infty) + \mathbf{b} \mid p \in \mathcal{L}_n\}$ 
4  $\mathbf{d} \leftarrow (-\infty, -\infty, \dots, -\infty)$  //  $m$ -dimensional
5 for  $k \in \{1, \dots, m\}$  do
6    $D_k \leftarrow \{\mathbf{d}'[k] \mid \mathbf{d}' \in D_p \cup D_n\}$ 
7    $\mathbf{d}[k] \leftarrow \arg \min_{d \in D_k} (|D_{FP}^k(d)| + |D_{FN}^k(d)|)$ , where
8    $D_{FP}^k(d) \leftarrow \{\mathbf{d}'[k] \mid \mathbf{d}'[k] > d, \mathbf{d}' \in D_n\}$ 
9    $D_{FN}^k(d) \leftarrow \{\mathbf{d}'[k] \mid \mathbf{d}'[k] \leq d, \mathbf{d}' \in D_p\}$ 
10 return  $d$ 
```

In Algorithm 16, m denotes the number of *within* operators and \mathbf{b} is the m -dimensional vector of deadlines associated with each *within* operator in the TWTL formula ϕ . We assume that the order of the *within* operators is given by the post-order traversal of $AST(\phi)$, i.e., recursively traversing the children nodes first and then the node itself.

Note that the optimization problem in line 7 of Algorithm 16 may have multiple optimal solutions, in which case we choose the lowest deadline. An example of this situation is shown in Sec. 10.4.

The complexity of the learning procedure is $O(2^{|\phi|+|AP|} + (|\mathcal{L}_p| + |\mathcal{L}_n|) \cdot l_m \cdot |\phi| + m^2 \cdot (|\mathcal{L}_p| + |\mathcal{L}_n|))$, where: (a) the first term is the complexity of constructing \mathcal{A}_∞ (line 1); (b) the second term corresponds to computing the tight deadlines for all traces positive and negative in lines 2 and 3, respectively; (c) the third term is the complexity of the for loop, which computes each deadline separately in a greedy fashion (lines 5–9). The maximum length of a trace (positive or negative) is denoted by l_m in the complexity formula.

9. TWTL Python package

We provide a Python 2.7 implementation named PyTWTL of the proposed algorithms based on LOMAP [42], ANTLRv3 [35] and NetworkX [15] libraries. PyTWTL implementation is released under the GPLv3 license and can be downloaded from hyness.bu.edu/twttl. The library can be used to:

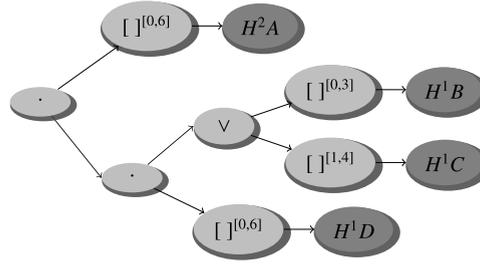


Fig. 2. The AST corresponding to the TWTL formula in Eq. (20).

1. construct a DFA \mathcal{A}_ϕ and an annotated DFA \mathcal{A}_∞ from a TWTL formula ϕ ;
2. monitor the satisfaction of a TWTL formula ϕ ;
3. monitor the satisfaction of an arbitrary relaxation of ϕ , i.e., $\phi(\infty)$;
4. compute the temporal relaxation of a trace with respect to a TWTL formula;
5. compute a satisfying control policy with respect to a TWTL formula ϕ ;
6. compute a minimally relaxed control policy with respect to a TWTL formula ϕ , i.e., for $\phi(\tau)$ such that $|\tau|_{TR}$ is minimal;
7. verify if all traces of a system satisfy some relaxed version of a TWTL formula ϕ ;
8. learn the parameters of a TWTL formula ϕ , i.e., the deadlines of the *within* operators in ϕ .

The parsing of TWTL formulae is performed using ANTLRv3 framework. We provide grammar files which may be used to generate lexers and parsers for other programming languages such as Java, C/C++, Ruby. To support Python 2.7, we used version 3.1.3 of ANTLRv3 and the corresponding Python runtime ANTLR library, which we included in our distribution for convenience.

10. Case studies

In this section, we present some examples highlighting the solutions for the verification, synthesis and learning problems. First, we show the automaton construction procedure on a TWTL formula and how the tight deadlines are inferred for a given trace. Then, we consider an example involving a robot whose motion is modeled as a TS. The policy computation algorithm is used to solve a path planning problem with rich specifications given as TWTL formulae. The procedure for performing verification, i.e., all robot trajectories satisfy a given TWTL specification, is also shown. Finally, the performance of the heuristic learning algorithm is demonstrated on a simple example.

10.1. Automata construction and temporal relaxation

Consider the following TWTL specification over the set of atomic propositions $AP = \{A, B, C, D\}$:

$$\phi = [H^2 A]^{0,6} \cdot ([H^1 B]^{0,3} \vee [H^1 C]^{1,4}) \cdot [H^1 D]^{0,6} \quad (20)$$

An AST of formula ϕ is shown in Fig. 2. The TWTL formula ϕ is converted to an annotated DFA \mathcal{A}_∞ using Algorithm 1. The procedure recursively constructs the DFA from the leaves of the AST to the root. A few processing steps are shown in Fig. 3. The construction of DFA corresponding to a leaf, i.e., a *hold* operator, is straightforward, see Fig. 3a. Next, the transformation corresponding to a *within* operator is shown in Fig. 3b. Note that the delay of one time unit is due to the lower bound of the time window of the *within* operator. Also, note that the automaton restarts on symbols that block the DFA corresponding to the inner formula $H^1 C$.

The next two figures, Fig. 3c and Fig. 3d, show the translation of the disjunction operator. Specifically, Fig. 3c, shows the product DFA corresponding to the disjunction without merging the final states. Since none of the final states have outgoing transitions, see Corollary 6.11, and they can be merged into a single final state, see Fig. 3d. However, we still need to keep track of which subformula of the disjunctions holds. The annotation variable $T.choices$, introduced in Sec. 7.3.2, stores this information as

$$\begin{cases} L = \{(s_{11}, B \wedge \neg C), (s_{11}, B \wedge C), (s_{12}, B \wedge \neg C)\}, \\ R = \{(s_{02}, \neg B \wedge C), (s_{02}, B \wedge C), (s_{12}, \neg B \wedge C)\}, \\ B = \{(s_{12}, B \wedge C)\}. \end{cases} \quad (21)$$

Notice that the tuples in Eq. (21) correspond to the ingoing edges of the final states in the DFA from Fig. 3c. Finally, the DFA corresponding to the overall specification formula ϕ is shown in Fig. 3e.

Let $\phi_A = [H^2 A]^{0,6}$, $\phi_B = [H^1 B]^{0,3}$, $\phi_C = [H^1 C]^{1,4}$, and $\phi_D = [H^1 D]^{0,6}$ be subformulae of ϕ associated with the *within* operators. The annotation data for these subformulae is shown in the following table.

Subformula	$T.I$	$T.F$
ϕ	$\{s_0\}$	$\{s_{10}\}$
ϕ_A	$\{s_0\}$	$\{s_3\}$
ϕ_B	$\{s_3, s_5, s_6\}$	$\{s_8\}$
ϕ_C	$\{s_3\}$	$\{s_3\}$
ϕ_D	$\{s_8\}$	$\{s_{10}\}$

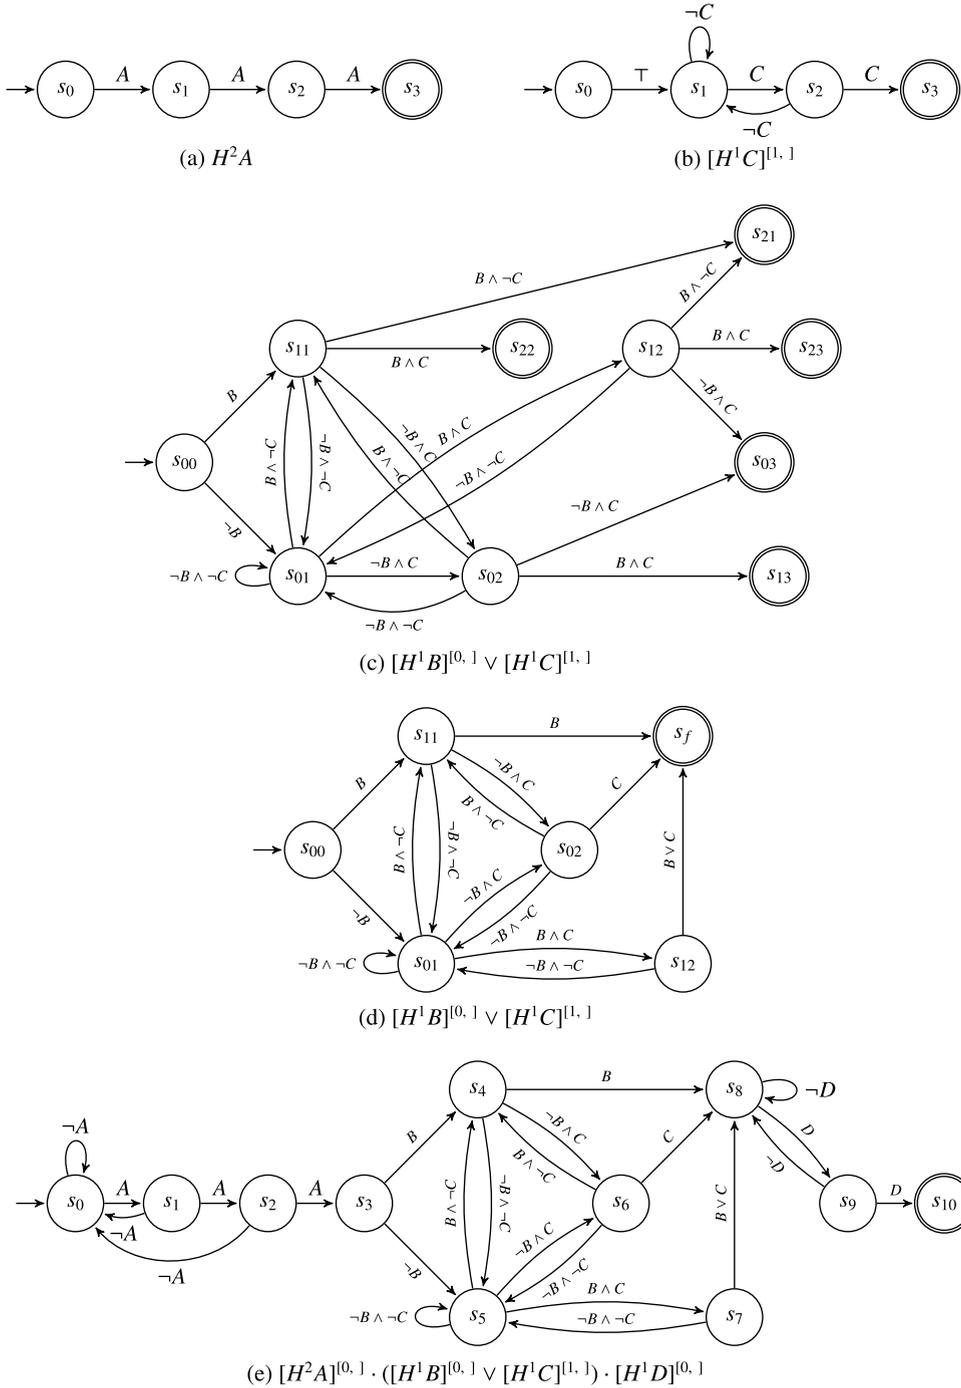


Fig. 3. Annotated automata corresponding to subformulae of the TWTL specification in Eq. (20).

Consider the following word over the alphabet $\Sigma = 2^{AP}$:

$$\sigma = \epsilon, \{A\}, \{A\}, \{A\}, \epsilon, \{B, C\}, \{B, C\}, \epsilon, \{D\}, \{D\} \quad (22)$$

where ϵ is the empty symbol. The following table shows the stages of [Algorithm 11](#) as the symbols of the word σ are processed:

No.	Symbol	State	ϕ_A	ϕ_B	ϕ_C	ϕ_D
Init		s_0	$(\top, \perp, 0)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$
0	ϵ	s_0	$(\top, \perp, 1)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$
1	$\{A\}$	s_1	$(\top, \perp, 2)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$
2	$\{A\}$	s_2	$(\top, \perp, 3)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$	$(\perp, \perp, -1)$
3	$\{A\}$	s_3	$(\perp, \top, 3)$	$(\top, \perp, 0)$	$(\top, \perp, 0)$	$(\perp, \perp, -1)$
4	ϵ	s_5	$(\perp, \top, 3)$	$(\top, \perp, 1)$	$(\top, \perp, 1)$	$(\perp, \perp, -1)$
5	$\{B, C\}$	s_7	$(\perp, \top, 3)$	$(\top, \perp, 2)$	$(\top, \perp, 2)$	$(\perp, \perp, -1)$
6	$\{B, C\}$	s_8	$(\perp, \top, 3)$	$(\perp, \top, 2)$	$(\perp, \top, 2)$	$(\top, \perp, 0)$
7	ϵ	s_8	$(\perp, \top, 3)$	$(\perp, \top, 2)$	$(\perp, \top, 2)$	$(\top, \perp, 1)$
8	$\{D\}$	s_9	$(\perp, \top, 3)$	$(\perp, \top, 2)$	$(\perp, \top, 2)$	$(\top, \perp, 2)$
9	$\{D\}$	s_{10}	$(\perp, \top, 3)$	$(\perp, \top, 2)$	$(\perp, \top, 2)$	$(\perp, \top, 2)$

where each 3-tuple in last four columns represents the annotation variables $T.ongoing$, $T.done$ and $T.steps$, respectively. The temporal relaxation for σ can be extracted from the values in the last row by subtracting the deadlines of the *within* operators from them. Thus, the vector of tightest τ values is $(-3, -1, -2, -3)$. However, because ϕ_B and ϕ_C are in disjunction, we have the temporal relaxation $\tau = (-3, -\infty, -2, -3)$, where we choose to ignore the subformula containing ϕ_B . Thus, the maximum temporal relaxation is $|\tau|_{TR} = -2$.

10.2. Control policy synthesis

Consider a robot moving in an environment represented as the finite graph shown in [Fig. 4a](#). The nodes of the graph represent the points of interest, while the edges indicate the possibility of moving the robot between the edges' endpoints. The numbers associated with the edges represent the travel times, and we assume that all the travel times are integer multiples of a time step Δt . The robot may also stay at any of the points of interest.

The motion of the robot is abstracted as a transition system \mathcal{T} , which is obtained from the finite graph by splitting each edge into a number of transitions equal to the corresponding edge's travel time. The generated transition system thus has 27 states and 67 transitions and is shown in [Fig. 4b](#).

Consider the TWTL specification ϕ from [Eq. \(20\)](#). The product automaton $\mathcal{P} = \mathcal{T} \times \mathcal{A}_\infty$ is constructed, where \mathcal{A}_∞ is the annotated DFA corresponding to $\phi(\infty)$ shown in [Fig. 3e](#). The product automaton \mathcal{P} has 204 states and 378 transitions. The control policy computed by using [Algorithm 14](#) is

$$\mathbf{x} = \text{Base}, A, A, A, C, C, \text{Base}, D, D, \quad (23)$$

which generates the output word

$$\sigma = \epsilon, \epsilon, \{A\}, \{A\}, \{A\}, \epsilon, \{C\}, \{C\}, \epsilon, \epsilon, \{D\}, \{D\}. \quad (24)$$

The minimum temporal relaxation for σ is $|\tau|_{TR} = -2$, where $\tau = (-2, -\infty, -2, -3)$ is the minimal temporal relaxation vector associated with σ .

10.3. Verification

In the verification problem, we are concerned with checking for the existence of relaxed specifications for every possible run of a transition system.

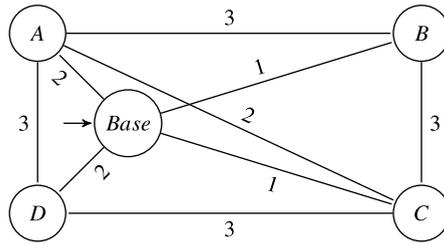
To illustrate this problem, consider the transition system in [Fig. 5](#) and the following two TWTL specifications:

$$\phi_1 = [H^1 A]^{[1,2]} \quad (25)$$

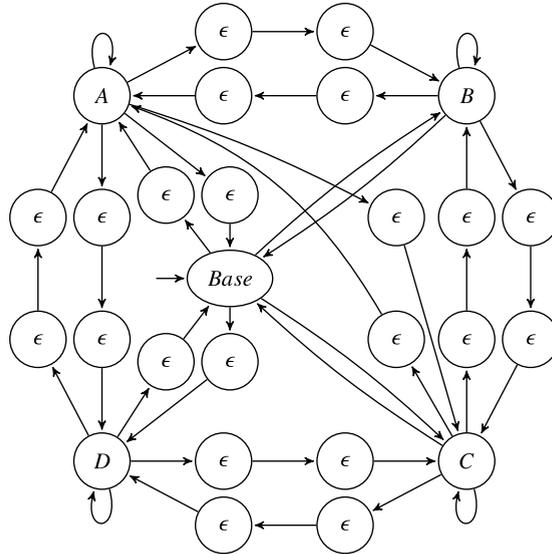
$$\phi_2 = [H^3 \neg B]^{[1,2]} \quad (26)$$

To check the transition system $\mathcal{T}^{\text{simple}}$ against the two specifications, we can use [Algorithm 15](#). It is straightforward that the procedure will return true for ϕ_1 , because every run of $\mathcal{T}^{\text{simple}}$ satisfies $\phi_1(3) = [H^1 A]^{[1,2+3]}$. Note that the runs of the transition system may not need to satisfy the original specification as the satisfaction of a relaxed version is sufficient. Similarly, [Algorithm 15](#) returns false for ϕ_2 , because there exists a run of $\mathcal{T}^{\text{simple}}$ that does not satisfy any relaxation of the specification, e.g., $\mathbf{x} = (A, B, B, \epsilon, A)^*$.

An important conclusion highlighted by the two examples is that the verification problem proposed in this paper is concerned with checking a system against the logical structure of a specification and not against any particular time bounds. This might be useful in situation where the deadlines of the specification are not known *a priori*, but the logical structure of the specification is.



(a) An environment with five points of interest, Base A, B, C, and D. The edges indicate the existence of paths between their endpoints, while the associated numbers represent the travel times of the edges. The robot may stay at a region of interest.



(b) The transition system \mathcal{T} obtained from the environment graph shown in Fig. 4a.

Fig. 4. The environment where the robot operates and its abstraction \mathcal{T} .

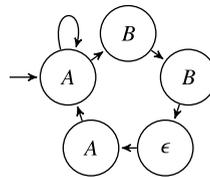


Fig. 5. A simple transition system \mathcal{T}^{simple} .

10.4. Learning deadlines from data

In the previous two cases, we use the TWTL specifications in conjunction with problems involving infinite sets of words encoded as transition systems. However, it is often the case that only finite sets of (output) trajectories are available. In this section, we give a simple example of the learning problem presented in Sec. 5.

Consider the specification $\phi_l = [H^1 A]^{[0, d_1]} \cdot [H^2 B]^{[0, d_2]}$ with unknown deadlines and the following set of labeled trajectories, where C_p and C_n are the positive and negative example labels, respectively:

Word	Label	Deadlines
$\sigma_1 = \{A\}, \{A\}, \{A\}, \{B\}, \{B\}, \{B\}, \{B\}, \epsilon$	C_p	(1, 3)
$\sigma_2 = \epsilon, \{A\}, \{A\}, \epsilon, \{B\}, \{B\}, \{B\}, \epsilon$	C_p	(2, 3)
$\sigma_3 = \{B\}, \epsilon, \{A\}, \{A\}, \{B\}, \{B\}, \{B\}, \{B\}$	C_n	(3, 2)
$\sigma_4 = \epsilon, \{A\}, \{A\}, \epsilon, \epsilon, \{B\}, \{B\}, \{B\}$	C_n	(2, 4)

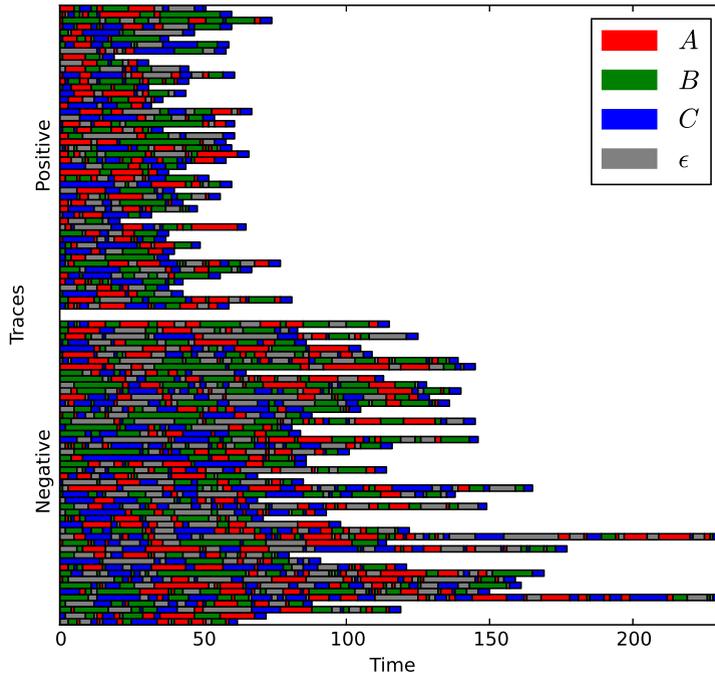


Fig. 6. The training set contains 50 positive and 50 negative labeled trajectories.

The last column in the above table shows the tight deadlines obtained in lines 2 and 3 of Algorithm 16. Next, the learning algorithm computes the heuristic sets D_{FP}^k and D_{FN}^k , $k \in \{d_1, d_2\}$, of false positive and false negative trajectories, respectively:

Deadline	Value	D_{FP}^k	D_{FN}^k	$ D_{FP}^k + D_{FN}^k $
d_1	1	\emptyset	$\{\sigma_2\}$	1
d_1	2	$\{\sigma_4\}$	\emptyset	1
d_1	3	$\{\sigma_3, \sigma_4\}$	\emptyset	2
d_2	2	$\{\sigma_3\}$	$\{\sigma_1, \sigma_2\}$	3
d_2	3	$\{\sigma_3\}$	\emptyset	1
d_2	4	$\{\sigma_3, \sigma_4\}$	\emptyset	2

Finally, Algorithm 16 chooses the deadline pair $\mathbf{d} = (d_1, d_2) = (1, 3)$ that has the lowest heuristic misclassification rate, $|D_{FP}^k| + |D_{FN}^k|$ shown in the last column of the above table, for d_1 and d_2 , respectively. An important observation is that the inferred formula $\phi_{\mathbf{d}} = [H^1 A]^{[0,1]} \cdot [H^2 B]^{[0,3]}$ has zero as actual misclassification rate. The discrepancy between the values in the table and the actual value of the final misclassification rate are due to the heuristic of synthesizing each deadline separately. Thus, the heuristic procedure in Algorithm 16 ignores the temporal and logical structure of the template TWTL formula which may lead to suboptimal performance, i.e., misclassification rate.

We also tested the learning algorithm on larger sets of trajectories. Algorithm 16 was ran using the template TWTL formula $[H^2 A]^{[0,d_1]} \cdot [H^3 B]^{[2,d_2]} \cdot [H^2 C]^{[0,d_3]}$. The inference was performed using a set of 100 trajectories, 50 positive and 50 negative, shown in Fig. 6. Executing Algorithm 16 returned the vector of deadlines $(d_1, d_2, d_3) = (29, 40, 31)$ that induces a misclassification rate of 14%.

11. Conclusion

In this paper, we introduced a specification language called time window temporal logic (TWTL), which is a linear-time logic encoding sets of discrete-time bounded-time trajectories. We showed that TWTL has several benefits over other bounded temporal logics in terms of complexity and easiness to express and comprehend specifications. Different from other temporal logics, TWTL has an explicit concatenation operator, which enables the compact representation of serial tasks. Such a compact representation significantly reduces the complexity of constructing the automaton for the accepting language. In this paper, we also presented temporal relaxations of TWTL formulae and provided provably-correct algorithms to construct an annotated automaton that can encode all temporal relaxations of a given TWTL formula. Moreover, we demonstrated the potential of TWTL and its relaxation on three problems related to verification, synthesis, and learning.

In the verification problem, we checked whether a system can satisfy the structure of a given formula without considering its time bounds. In the synthesis problem, we found a control policy for a system that satisfies the original TWTL formula or its minimal relaxation in case of infeasibility. In the learning problem, we considered a data set and a template TWTL formula with parametric time bounds, and we synthesized the time parameters by minimizing the misclassification rate. Finally, we developed a Python package for the solutions of the aforementioned problems.

As future work, we plan to extend the application areas of the proposed methods. For example, TWTL is a good fit for statistical model checking, where the problem of learning deadlines can be modified to yield statistically robust deadline values in a template formula. The current version of the learning algorithm can find the time bounds of a given template formula (with fixed structure) from a data set. We are also working on more advanced algorithms that can infer not only the time bounds but also the structure of the template. Furthermore, we plan to improve the Python package PyTWTL by integrating automata minimization in the construction procedure in a way that (i) preserves annotation, and (ii) takes into account structure of the generated automata. Since the languages associated with TWTL formulae are finite, specialized minimization techniques may be used. For instance, one approach is to use Deterministic Finite Cover Automata [24,25,8] that can decrease the return automata's sizes significantly with almost no additional computational cost. Other small optimizations we plan to include in PyTWTL are: (i) the case when the satisfaction of atomic propositions is assumed mutually exclusive; and (ii) preprocessing of TWTL formulae for performance improvement using AST rewriting rules. We also plan to develop AST rewriting rules to automatically transform a TWTL formula to DFW form.

References

- [1] Derya Aksaray, Kevin Leahy, Calin Belta, Distributed multi-agent persistent surveillance under temporal logic constraints, in: 5th IFAC Workshop on Distributed Estimation and Control in Networked Systems, NecSys, IFAC-PapersOnLine 48 (22) (2015) 174–179, <http://dx.doi.org/10.1016/j.ifacol.2015.10.326>, URL <http://www.sciencedirect.com/science/article/pii/S240589631502217X>.
- [2] Derya Aksaray, Cristian-Ioan Vasile, Calin Belta, Dynamic routing of energy-aware vehicles with temporal logic constraints, in: IEEE International Conference on Robotics and Automation, ICRA, May 2016, pp. 3141–3146.
- [3] Rajeev Alur, Kousha Etesami, Salvatore La Torre, Doron Peled, Parametric temporal logic for “model measuring”, ACM Trans. Comput. Log. (ISSN 1529-3785) 2 (3) (July 2001) 388–407, <http://dx.doi.org/10.1145/377978.377990>.
- [4] Eugene Asarin, Alexandre Donzé, Oded Maler, Dejan Nickovic, Parametric identification of temporal properties, in: Runtime Verification: Second International Conference, RV, Springer, Berlin, Heidelberg, ISBN 978-3-642-29860-8, 2012, pp. 147–160.
- [5] Christel Baier, Joost-Pieter Katoen, Principles of Model Checking, MIT Press, ISBN 978-0-262-02649-9, 2008.
- [6] Calin Belta, Volkan Isler, George J. Pappas, Discrete abstractions for robot planning and control in polygonal environments, IEEE Trans. Robot. 21 (5) (2005) 864–874, <http://dx.doi.org/10.1109/TRO.2005.851359>.
- [7] Cezar Cămpăanu, K. Culik II, Kai Salomaa, Sheng Yu, State complexity of basic operations on finite languages, in: Oliver Boldt, Helmut Jürgensen (Eds.), Automata Implementation, in: Lecture Notes in Comput. Sci., vol. 2214, Springer, Berlin, Heidelberg, ISBN 978-3-540-42812-1, 2001, pp. 60–70.
- [8] Cezar Cămpăanu, Andrei Păun, Jason R. Smith, Incremental construction of minimal deterministic finite cover automata, Theoret. Comput. Sci. 363 (2) (2006) 135–148, <http://dx.doi.org/10.1016/j.tcs.2006.07.020>, URL <http://www.sciencedirect.com/science/article/pii/S0304397506004567>.
- [9] Jan Daciuk, Comparison of construction algorithms for minimal, acyclic, deterministic, finite-state automata from sets of strings, in: Implementation and Application of Automata, Springer, 2003, pp. 255–261, URL <http://dl.acm.org/citation.cfm?id=1756384.1756412>.
- [10] Alexandre Duret-Lutz, Manipulating LTL formulas using Spot 1.0, in: Proceedings of the 11th International Symposium on Automated Technology for Verification and Analysis, ATVA, in: Lecture Notes in Comput. Sci., vol. 8172, Springer, Hanoi, Vietnam, October 2013, pp. 442–445.
- [11] Georgios E. Fainekos, Antoine Girard, Hadas Kress-Gazit, George J. Pappas, Temporal logic motion planning for dynamic robots, Automatica 45 (2) (2009) 343–352, <http://dx.doi.org/10.1016/j.automatica.2008.08.008>, URL <http://www.sciencedirect.com/science/article/pii/S000510980800455X>.
- [12] Yuan Gao, Kai Salomaa, Sheng Yu, Transition complexity of incomplete DFAs, Fund. Inform. (ISSN 0169-2968) 110 (1–4) (January 2011) 143–158, URL <http://dl.acm.org/citation.cfm?id=2362097.2362107>.
- [13] Bombarda Giuseppe, Vasile Cristian Ioan, Penedo Alvarez Francisco, Hirotohi Yasuoka, Belta Calin, A decision tree approach to data classification using signal temporal logic, in: Hybrid Systems: Computation and Control, HSCC, Vienna, Austria, April 2016, pp. 1–10.
- [14] Meng Guo, Dimos V. Dimarogonas, Multi-agent plan reconfiguration under local LTL specifications, Int. J. Robot. Res. 34 (2) (2015) 218–235, <http://dx.doi.org/10.1177/0278364914546174>, URL <http://ijr.sagepub.com/content/34/2/218.abstract>.
- [15] Aric A. Hagberg, Daniel A. Schult, Pieter J. Swart, Exploring network structure, dynamics, and function using NetworkX, in: Proceedings of the 7th Python in Science Conference, CA, USA, August 2008, pp. 11–15.
- [16] Yo-Sub Han, Kai Salomaa, State complexity of union and intersection of finite languages, in: Tero Harju, Juhani Karhumäki, Arto Lepistö (Eds.), Developments in Language Theory, in: Lecture Notes in Comput. Sci., vol. 4588, Springer, Berlin, Heidelberg, ISBN 978-3-540-73207-5, 2007, pp. 217–228.
- [17] John E. Hopcroft, Rajeev Motwani, Jeffrey D. Ullman, Introduction to Automata Theory, Languages, and Computation, 3rd edition, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, ISBN 0321455363, 2006.
- [18] Sumit K. Jha, Edmund M. Clarke, Christopher J. Langmead, Axel Legay, André Platzer, Paolo Zuliani, A Bayesian approach to model checking biological systems, in: Proceedings of the 7th Int. Conference on Computational Methods in Systems Biology, CMSB'09, Springer-Verlag, Berlin, Heidelberg, ISBN 978-3-642-03844-0, 2009, pp. 218–234.
- [19] Xiaoqing Jin, Alexandre Donzé, Jyotirmoy V. Deshmukh, Sanjit A. Seshia, Mining requirements from closed-loop control models, IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. (ISSN 0278-0070) 34 (11) (Nov 2015) 1704–1717, <http://dx.doi.org/10.1109/TCAD.2015.2421907>.
- [20] Sertac Karaman, Emilio Frazzoli, Vehicle routing problem with metric temporal logic specifications, in: IEEE Conference on Decision and Control, CDC, December 2008, pp. 3953–3958.
- [21] Kangjin Kim, Georgios Fainekos, Sriram Sankaranarayanan, On the minimal revision problem of specification automata, Int. J. Robot. Res. 34 (12) (2015) 1515–1535, <http://dx.doi.org/10.1177/0278364915587034>, URL <http://ijr.sagepub.com/content/early/2015/08/11/0278364915587034.abstract>.
- [22] Marius Kloetzer, Calin Belta, A fully automated framework for control of linear systems from temporal logic specifications, IEEE Trans. Automat. Control 53 (1) (2008) 287–297, <http://dx.doi.org/10.1109/TAC.2007.914952>.
- [23] Zhaodan Kong, Austin Jones, Ana Medina Ayala, Ebru Aydın Gol, Calin Belta, Temporal logic inference for classification and prediction from data, in: Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control, HSCC'14, ACM, New York, NY, USA, ISBN 978-1-4503-2732-9, 2014, pp. 273–282.
- [24] Heiko Körner, On minimizing cover automata for finite languages in $O(N \log N)$ time, in: Proceedings of the 7th International Conference on Implementation and Application of Automata, CIAA'02, Springer-Verlag, Berlin, Heidelberg, ISBN 3-540-40391-4, 2003, pp. 117–127.

- [25] Heiko Körner, A time and space efficient algorithm for minimizing cover automata for finite languages, *Internat. J. Found. Comput. Sci.* 14 (6) (2003) 1071–1086, <http://dx.doi.org/10.1142/S0129054103002187>.
- [26] Ron Koymans, Specifying real-time properties with metric temporal logic, *Real-Time Syst.* 2 (4) (1990) 255–299, <http://dx.doi.org/10.1007/BF01995674>.
- [27] Hadas Kress-Gazit, Georgios E. Fainekos, George J. Pappas, Temporal-logic-based reactive mission and motion planning, *IEEE Trans. Robot.* 25 (6) (2009) 1370–1381, <http://dx.doi.org/10.1109/TRO.2009.2030225>.
- [28] Orna Kupferman, Moshe Y. Vardi, Model checking of safety properties, *Form. Methods Syst. Des.* 19 (3) (2001) 291–314, <http://dx.doi.org/10.1023/A:1011254632723>.
- [29] Timo Latvala, Efficient model checking of safety properties, in: *10th International SPIN Workshop, Model Checking Software*, Springer, 2003, pp. 74–88.
- [30] Kevin Leahy, Dingjiang Zhou, Cristian-Ioan Vasile, Konstantinos Oikonomopoulos, Mac Schwager, Calin Belta, Provably correct persistent surveillance for unmanned aerial vehicles subject to charging constraints, in: *International Symposium on Experimental Robotics*, ISER, Marrakech/Essaouira, Morocco, June 2014.
- [31] Scott C. Livingston, Pavithra Prabhakar, Alex B. Jose, Richard M. Murray, Patching task-level robot controllers based on a local μ -calculus formula, in: *International Conference on Robotics and Automation*, ICRA, 2013, pp. 4588–4595.
- [32] Eva Maia, Nelma Moreira, Rogério Reis, Incomplete transition complexity of some basic operations, in: Peter van Emde Boas, Frans C.A. Groen, Giuseppe F. Italiano, Jerzy Nawrocki, Harald Sack (Eds.), *SOFSEM 2013: Theory and Practice of Computer Science*, in: *Lecture Notes in Comput. Sci.*, vol. 7741, Springer, Berlin, Heidelberg, ISBN 978-3-642-35842-5, 2013, pp. 319–331.
- [33] Oded Maler, Dejan Nickovic, Monitoring temporal properties of continuous signals, in: Yassine Lakhnech, Sergio Yovine (Eds.), *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Springer, 2004, pp. 152–166.
- [34] Zohar Manna, Amir Pnueli, Verification of Concurrent Programs. Part I. The Temporal Framework, Technical report, DTIC Document, 1981.
- [35] Terence Parr, *The Definitive ANTLR Reference: Building Domain-Specific Languages*, Pragmatic Bookshelf, ISBN 978-0978739256, 2007.
- [36] Marco Pavone, Nabendra Bisnik, Emilio Frazzoli, Volkan Isler, A stochastic and dynamic vehicle routing problem with time windows and customer impatience, *Mob. Netw. Appl.* 14 (3) (2009) 350–364, <http://dx.doi.org/10.1007/s11036-008-0101-1>.
- [37] Luis I. Reyes Castro, Pratik Chaudhari, Jana Tumova, Sertac Karaman, Emilio Frazzoli, Daniela Rus, Incremental sampling-based algorithm for minimum-violation motion planning, in: *IEEE Conference on Decision and Control*, CDC, December 2013, pp. 3217–3224.
- [38] Marius M. Solomon, Algorithms for the vehicle routing and scheduling problems with time window constraints, *Oper. Res.* 35 (2) (1987) 254–265, <http://dx.doi.org/10.1287/opre.35.2.254>.
- [39] Ilya Tkachev, Alessandro Abate, Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems, in: *Proceedings of the 16th Int. Conference on Hybrid Systems: Computation and Control*, Philadelphia, PA, April 2013, pp. 283–292.
- [40] Jana Tumova, Gavin C. Hall, Sertac Karaman, Emilio Frazzoli, Daniela Rus, Least-violating control strategy synthesis with safety rules, in: *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, HSCC'13, ACM, New York, NY, USA, ISBN 978-1-4503-1567-8, 2013, pp. 1–10.
- [41] Jana Tumova, Alejandro Marzinotto, Dimos V. Dimarogonas, Danica Kragic, Maximally satisfying LTL action planning, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, IROS, September 2014, pp. 1503–1510.
- [42] Alphan Ulusoy, Stephen L. Smith, Xu Chu Ding, Calin Belta, Daniela Rus, Optimality and robustness in multi-robot path planning with temporal logic constraints, *Int. J. Robot. Res.* 32 (8) (2013) 889–911, <http://dx.doi.org/10.1177/0278364913487931>.
- [43] Cristian-Ioan Vasile, Calin Belta, An automata-theoretic approach to the vehicle routing problem, in: *Proceedings of the Robotics: Science and Systems*, RSS, Berkeley, California, USA, July 2014, <http://dx.doi.org/10.15607/RSS.2014.X.045>.
- [44] Tichakorn Wongpiromsarn, Ufuk Topcu, Richard M. Murray, Receding horizon temporal logic planning for dynamical systems, in: *IEEE Conference on Decision and Control*, CDC, 2009, pp. 5997–6004.
- [45] Tichakorn Wongpiromsarn, Ufuk Topcu, Richard M. Murray, Receding horizon control for temporal logic specifications, in: *Proceedings of the 13th Int. Conference on Hybrid Systems: Computation and Control*, ACM, 2010, pp. 101–110.
- [46] Hengyi Yang, Bardh Hoxha, Georgios Fainekos, Querying parametric temporal logic properties on embedded systems, in: *24th IFIP WG 6.1 International Conference Testing Software and Systems*, ICTSS, Springer, Berlin, Heidelberg, ISBN 978-3-642-34691-0, 2012, pp. 136–151.