

# Safe Planning through Incremental Decomposition of Signal Temporal Logic Specifications

Parv Kapoor<sup>1</sup>, Eunsuk Kang<sup>1</sup>, and Rômulo Meira-Góes<sup>2</sup>

<sup>1</sup> Carnegie Mellon University, Pittsburgh, PA, USA

<sup>2</sup> Pennsylvania State University, State College, PA, USA

parvk@cs.cmu.edu, eunsukk@andrew.cmu.edu, romulo@psu.edu

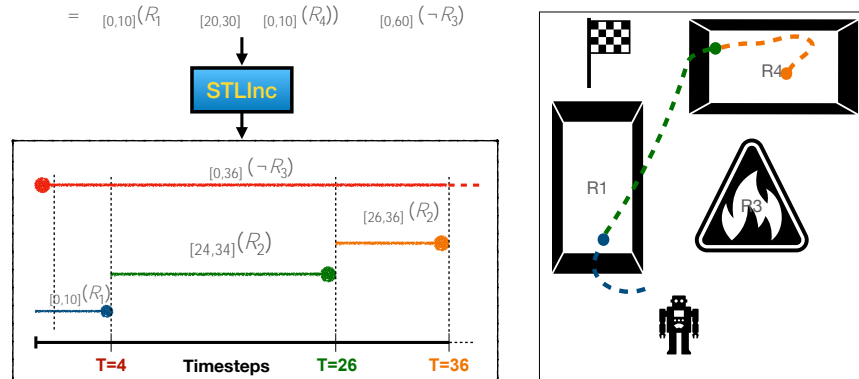
**Abstract.** *Trajectory planning* is a critical process that enables autonomous systems to safely navigate complex environments. *Signal temporal logic (STL)* specifications are an effective way to encode complex, temporally extended objectives for trajectory planning in cyber-physical systems (CPS). However, the complexity of planning with STL using existing techniques scales exponentially with the number of nested operators and the time horizon of a given specification. Additionally, poor performance is exacerbated at runtime due to limited computational budgets and compounding modeling errors. Decomposing a complex specification into smaller subtasks and incrementally planning for them can remedy these issues. In this work, we present a method for decomposing STL specifications to improve planning efficiency and performance. The key insight in our work is to encode all specifications as a set of basic constraints called *reachability* and *invariance constraints*, and schedule these constraints sequentially at runtime. Our experiment shows that the proposed technique outperforms the state-of-the-art trajectory planning techniques for both linear and non-linear dynamical systems.

**Keywords:** Signal Temporal Logic · Planning · Cyber Physical Systems

## 1 Introduction

Most autonomous robots interacting with the physical world need to achieve complex objectives while dealing with uncertainty and stochasticity in their environment. This problem is exacerbated by short response times expected while ensuring runtime efficiency. Hence, formulating these complex objectives accurately is a crucial step in realizing the desired behaviors for robotic operations.

Temporal logics such as *linear temporal logic (LTL)* [19] and *signal temporal logic (STL)* [16] provide a precise way to encode objectives that are expressed in a natural language. STL has received special attention in the community due to its rich quantitative semantics that can quantitatively measure satisfaction of a given property that encodes an objective. Additionally, it can be used to describe complex properties over real valued signals such as state trajectories arising from continuous dynamical systems. For robotic planning, STL can be used to



**Fig. 1.** **Left:** An STL specification with multiple nested temporal operators and a possible decomposition into subtasks. **Right:** A sample trajectory that satisfies the specification in a planar environment.

describe complex behaviors with concrete time deadlines such as those found in trajectory planning and task planning. Planners can use these specifications to generate specification-conforming behavior.

A significant amount of common robotic objectives can be interpreted as a sequence of subtasks. It has been shown that incremental subtask planning can be done more efficiently compared to planning for a composite task [5, 7, 18]. However, when STL is used to represent these composite tasks, incremental planning becomes challenging. This issue is because STL semantics can encode the sequential nature of tasks but does not expose this structure to the planner. In such cases, the planners are forced to work with complex long-horizon specifications. When the horizon of the specification is longer than the planning horizon, planners can often generate suboptimal or violating plans. This problem is exacerbated when planning occurs at runtime with computational constraints and compounding modeling errors [2].

In this work, we propose a theory to decompose long-horizon, arbitrarily nested specifications into sub-specifications that can be satisfied incrementally. We define recursive rules for decomposition and propose a novel scheduling algorithm for incremental task planning. The key insight here is to “divide and conquer” STL requirements while ensuring, by construction, that the resulting plan satisfies the original composite specification. We illustrate the effectiveness of our proposed approach over an experiment involving robot exploration problems with linear and non-linear dynamics. Our preliminary experiment shows that our approach is able to more efficiently generate plans for complex, composite specifications in comparison to the existing state-of-the-art STL-based planning methods. In addition, our decomposition technique is agnostic to the

underlying system dynamics and the choice of planner, and can potentially be adapted by different planners.

The key contributions of this paper are:

- { A method for decomposing an STL specification into a set of smaller STL specifications that represent subtasks (Section 4.2);
- { A planning algorithm that incrementally schedules and executes these subtasks (Sections 4.3 and 4.4);
- { An evaluation of the proposed approach over a benchmark of motion planning tasks (Section 5).

## 2 Motivation

We illustrate the problem of planning from complex specifications using an example from the motion planning domain. We use a planning problem similar to the one defined in [14].

As illustrated in Figure 1, the goal of the agent (robot) is to visit regions  $R_1$  and  $R_4$  sequentially while avoiding an unsafe region,  $R_3$ . Additionally, upon reaching  $R_4$ , the agent needs to stay in it for 10 time steps. We combine three common motion planning patterns such as sequenced visit, stabilization, and global avoidance to create a specification with timed deadlines as follows:

$$\begin{aligned} &= \varphi_1 \wedge \varphi_2 \\ \varphi_1 &= \Box_{[0:10]}(R_1 \wedge \Box_{[20:30]} \Box_{[0:10]}(R_4)) \\ \varphi_2 &= \Box_{[0:60]}(\neg R_3) \end{aligned}$$

The state-of-the-art (SOTA) technique for planning from  $\varphi$ , originally proposed in [20], involves encoding the STL specification and the system dynamics as Mixed Integer Program (MIP) constraints and solving the constrained optimization problem in a receding horizon fashion. A new binary decision variable is introduced for each atomic proposition per time step in the STL specification. A known drawback of this technique is its exponential worst-case complexity with respect to the number of binary variables [12]. Various encoding modifications have been suggested to enhance the efficiency of the technique by reducing the number of variables and constraints [21, 12].

However, even with reduced variable encoding, current methods excel primarily with short-horizon specifications. When encoding nested temporal operators, a large number of additional variables and constraints are needed to capture the relationship between different temporal operators, in contrast to non-nested operators, where temporal constraints associated with each operator are considered independently. For example, let us take the subformula  $\Box_{[20:30]} \Box_{[0:10]}(R_4)$  from  $\varphi_1$ . For encoding this subformula into an MIP, we would need 11 (outer eventually) + 121 (inner always) = 132 binary variables.<sup>3</sup> In general, as the nesting depth increases, the number of variables can increase exponentially.

<sup>3</sup> For the outer  $\Box$  clause, 11 binary variables are introduced to encode that the inner clause is satisfied within interval [20,30]; for each time point in [0,10] interval, another set of 11 variables are introduced, thus resulting in  $11 \cdot 11 = 121$  variables.

In this work, we propose a technique to improve the scalability of STL planning algorithms through decomposition of STL specifications. Our idea is inspired by human planning, where long-term goals are achieved by breaking tasks into incremental sub-goals [8]. Concretely, by decomposing the specification, we effectively remove the complexity of nested operators and also reduce the length of the lookahead horizon.

A possible decomposition of the specification into four subtasks is as follows:

$$\begin{aligned} sch_1 &= [0;10](R_1) \wedge [0;10](\neg R_3) \\ sch_2 &= [t_{R_1}+20;t_{R_1}+30](R_4) \wedge [t_{R_1}+20;t_{R_1}+30](\neg R_3) \\ sch_3 &= [t_{R_4}+0;t_{R_4}+10](\neg R_3 \wedge R_4) \\ sch_4 &= [t_{R_4},60](\neg R_3) \end{aligned}$$

Here, the symbolic time variables ( $t_{R_1}; t_{R_4}; t_{R_4,2}$ ) indicate when those subtasks get satisfied. More specifically,  $t_{R_4}$  indicates when the agent reaches Region 1 and  $t_{R_4,2}$  indicates when the agent has been inside Region 4 for 10 timesteps after reaching it. These variables then shift the time intervals of the other constraints that depend on them (e.g., time  $t_{R_1}$  from  $sch_1$  is used to concretize the time intervals for  $sch_2$ , whose time of satisfaction, in turn, influences  $sch_3$ ). These subtasks have shorter time horizons and no nested temporal operators, resulting in MIP constraints that are less complex than those that would result from composite specifications. As shown later in Section 5, this decomposition-based approach has potential to significantly improve the efficiency of planning.

### 3 Preliminaries

STL is a logical formalism used to define properties of continuous time real valued signals [16]. A signal  $\mathbf{s}$  is a function  $\mathbf{s} : \mathbb{T} \rightarrow \mathbb{R}^n$  that maps a time domain  $T \subseteq \mathbb{R}_{\geq 0}$  to a real valued vector. Then, an STL formula is defined as:

$$\varphi ::= j \mid \neg j \mid j \wedge j \mid j \vee j \mid U_{[a;b]}$$

where  $j$  is a predicate on the signal  $\mathbf{s}$  at time  $t$  in the form of  $j(\mathbf{s}(t)) > 0$  and  $[a;b]$  is the time interval (or simply  $I$ ). The *until* operator  $U$  defines that  $j$  must be true until  $k$  becomes true within a time interval  $[a;b]$ . Two other operators can be derived from *until*: *eventually* ( $F_{[a;b]} \varphi := \exists U_{[a;b]} \varphi$ ) and *always* ( $G_{[a;b]} \varphi := \forall F_{[a;b]} \varphi$ ).

**Definition 1.** Given a signal  $S_t$  representing a signal starting at time  $t$ , the Boolean semantics of satisfaction of  $S_t \models \varphi$  are defined inductively as follows:

$$\begin{aligned} S_t \models \varphi & \quad (\varphi) \quad (S(t)) > 0 \\ S_t \models \neg \varphi & \quad (\varphi) \quad \neg (S_t \models \varphi) \\ S_t \models \varphi_1 \wedge \varphi_2 & \quad (\varphi) \quad (S_t \models \varphi_1) \wedge (S_t \models \varphi_2) \\ S_t \models F_{[a;b]}(\varphi) & \quad (\varphi) \quad \exists t^\theta \in [t+a; t+b] \text{ s.t. } S_{t^\theta} \models \varphi \\ S_t \models G_{[a;b]}(\varphi) & \quad (\varphi) \quad \forall t^\theta \in [t+a; t+b] \text{ s.t. } S_{t^\theta} \models \varphi \end{aligned}$$

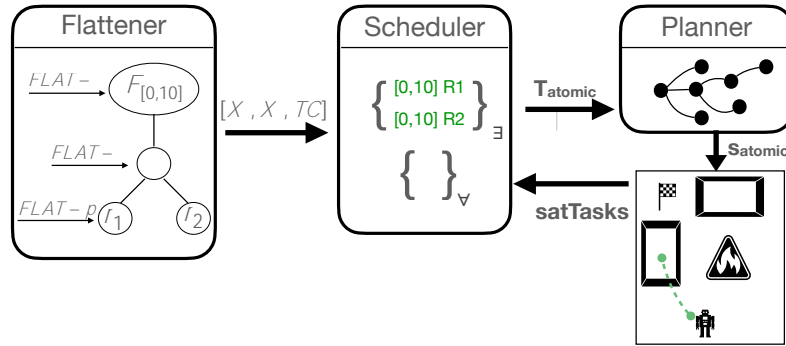


Fig. 2. Overview of the StlInc approach

Apart from the Boolean semantics, quantitative semantics are defined for a signal to compute a real-valued metric indicating *robustness*, i.e., the strength of satisfaction or violation.

**Definition 2.** Given a signal  $S_t$  representing a signal starting at time  $t$ , the quantitative semantics of satisfaction of  $S_t \models$  are defined inductively as follows:

$$\begin{aligned}
 (S_t; c) &= (x_t) \quad c \\
 (S_t; \neg) &= \neg (S_t; \cdot) \\
 (S_t; \neg_1 \wedge \neg_2) &= \min((S_t; \neg_1); (S_t; \neg_2)) \\
 (S_t; F_{[a;b]}(\cdot)) &= \max_{t' \in [t+a; t+b]} (S_{t'}^o; \cdot) \\
 (S_t; G_{[a;b]}(\cdot)) &= \min_{t' \in [t+a; t+b]} (S_{t'}^o; \cdot)
 \end{aligned}$$

For example, suppose that we are given (1)  $G_{[0;3]}(\text{distToR3}(t) \geq 3.0)$ , which states that the agent should maintain at least 3.0 meters away from region  $R_3$  for the next 4 time steps and (2) signal  $S_t$  that contains sequence  $\{3.0; 2.5; 3.0; 3.5\}$  for  $\text{distToR3}$ . Evaluating the robustness of satisfaction of  $\neg$  over  $S_t$  would result in a value of 0.5, implying that the agent violates the property by a degree of 0.5 (i.e., it fails to stay away from  $R_3$  by 0.5 meters).

## 4 Approach

### 4.1 Basic Concepts and Definitions

The overview of our planning framework (StlInc) is shown in Figure 2. The key idea behind this approach is that a bounded STL formula in our fragment can be decomposed into a finite set of the following two types of *task constraints*, each of which is associated with time interval  $I = [a; b]$  and state proposition  $p$ :

*Reachability*: The system ensures that  $\rho$  holds over at least one time step  $t$  within  $I$ .

*Invariance*: The system ensures that  $\rho$  holds over every step  $t$  within  $I$ .

Based on this idea, the framework carries out the incremental planning process over three steps. First, the *flattener* takes a user-specified STL specification ( ) and decomposes it into two sets of task constraints,  $X^{\rho}$  and  $X^{\beta}$ , which contain the reachability and invariance constraints, respectively. The decomposition is performed such that satisfying all of the constraints in these two sets implies the satisfaction of the original formula .

Next, the *scheduler* takes the two sets,  $X^{\rho}$  and  $X^{\beta}$ , and generates a sequence of *atomic tasks*,  $\tau = \{at_0; at_1; \dots; at_k\}$ , where (1) each atomic task  $at$  is a non-nested STL formula consisting of  $G_{[a;b]}(\rho)$  or  $F_{[a;b]}(\rho)$  (where  $\rho$  is a propositional formula) and (2) two different atomic tasks do not overlap in their time intervals. After this sequence is generated, the *planner* executes these atomic tasks one by one in the designated order. Once all tasks are executed, the system will have fulfilled  $X^{\rho}$  and  $X^{\beta}$ , thus satisfying the original goal of . The rest of this section describes each of the three steps in detail.

**STL fragment** Our approach is designed to handle specifications written in the following fragment of STL:

$$\begin{aligned} & ::= F_{[a;b]}(\ ) \ j \ G_{[a;b]}(\ ) \ j \ \_1 \ \wedge \ \_2 \ j \ \rho \\ \rho & ::= \rho_1 \ \wedge \ \rho_2 \ j \ \rho_1 \ \_ \ \rho_2 \ j \ : \ \rho \ j \end{aligned}$$

where  $\rho$  is a propositional formula that does not contain any temporal operator, and  $\_i$  is an atomic proposition. We target this STL fragment as (1) it significantly simplifies the decomposition process and (2) it is still expressive enough to capture many common behavioral patterns in the robotics planning domain [17]. Note that the fragment allows a nesting of temporal operators of an arbitrary depth, which is important for specifying sequential tasks. For example, the objective of *visiting* locations in a particular order can be defined as:

$$F_{[1;5]}(l_1 \ \wedge \ F_{[1;10]}(l_2 \ \wedge \ F_{[1;7]}(l_3)))$$

where  $l_1; l_2; l_3$  represent the locations to be visited.

**Task constraints** As mentioned at the beginning of the section, an STL formula imposes two types of constraints over system behavior: *reachability* and *invariance* constraints, which are formally defined as follows:

**De nition 3 (Reachability set)**. A *reachability set*,  $X^{\rho} \subseteq \mathcal{P}(T \times T \times U)$  is a set of tuples of form  $(l; h; \rho)$ , each stating that there exists some time  $t$  in interval  $[l; h]$ , inclusively, such that proposition  $\rho$  holds over the system state at time  $t$ .

**De nition 4 (Invariance set)**. An *invariance set*,  $X^{\beta} \subseteq \mathcal{P}(T \times T \times U)$ , is a set of tuples of form  $(l; h; \rho)$ , each stating that for every time  $t$  in interval  $[l; h]$ , inclusively, proposition  $\rho$  holds over the system state at time  $t$ .

$$\begin{aligned}
 & \frac{= p}{X^\circ = f(0;0;p)g \quad X^\beta = ; \quad TC = ;} \text{Flat-}p \quad \frac{\begin{array}{l} (X_1^\circ; X_1^\beta; TC_1) = \text{Flat}(\ \_1) \\ (X_2^\circ; X_2^\beta; TC_2) = \text{Flat}(\ \_2) \end{array}}{X^\circ = X_1^\circ \wedge X_2^\circ \\ X^\beta = X_1^\beta \wedge X_2^\beta \\ TC = TC_1 \wedge TC_2} \text{Flat-}\wedge \\
 & \frac{= F_{[a;b] \ \_1} \quad (X_1^\circ; X_1^\beta; TC_1) = \text{Flat}(\ \_1)}{X^\circ = \bar{f}(t+l; t+h; p) \ j \ \mathcal{Q}(l; h; p) \ \geq X_1^\circ g \\ X^\beta = \bar{f}(t+l; t+h; p) \ j \ \mathcal{Q}(l; h; p) \ \geq X_1^\beta g \\ TC = TC_1 \wedge \bar{f}(a; b; t)g} \text{Flat-}9 \\
 & \frac{= G_{[a;b] \ \_1} \quad (X_1^\circ; X_1^\beta; TC_1) = \text{Flat}(\ \_1)}{X^\circ = \bar{f}(k+l; k+l; p) \ j \ \mathcal{Q}(l; h; p) \ \geq X_1^\circ \wedge k \ \geq [a; b]g \\ X^\beta = \bar{f}(a+l; b+h; p) \ j \ \mathcal{Q}(l; h; p) \ \geq X_1^\beta g \\ TC = TC_1} \text{Flat-}8
 \end{aligned}$$

**Fig. 3.** Flattening rules

We introduce an additional set of constraints that specify intervals over symbolic time variables that are introduced during the flattening of  $F$  formulas:

**Definition 5 (Time variable intervals).** A time variable interval set,  $TC \subseteq \mathcal{P}(\mathbb{T} \times \mathbb{T} \times \mathbb{T})$ , is a set of tuples of the form  $(l, h, v)$ , each stating that symbolic time variable  $v$  takes a value in the interval  $[l; h]$ , inclusively.

## 4.2 Flattening

Given an input STL specification,  $\varphi$ , the goal of *flattening* is to construct two sets of constraints—reachability and invariance sets—whose satisfaction also implies the satisfaction of  $\varphi$ . Flattening is applied recursively based on the structure of  $\varphi$ , as shown through the rules in Figure 3. Along with  $X^\beta$  and  $X^\circ$ , each recursive step produces an additional auxiliary output  $TC$ , which is later used by the scheduler to resolve symbolic time variables.

**Flat- $p$**  In the basic case where  $\varphi = p$ , flattening generates one reachability constraint that requires  $p$  to be satisfied at the current time (i.e.,  $l = h = 0$ ). Hence, the reachability set for  $\varphi$  is  $(0;0;p)$ .

**Flat- $\wedge$**  Given conjunctive formula  $\varphi = \varphi_1 \wedge \varphi_2$ , the invariance set for  $\varphi$  is the union of  $X_1^\beta$  and  $X_2^\beta$ ; i.e., every invariance constraint in  $\varphi_1$  and  $\varphi_2$  must be satisfied. Similarly, the reachability set for  $\varphi$  is the union of  $X_1^\circ$  and  $X_2^\circ$ ; i.e., every reachability constraint in  $\varphi_1$  and  $\varphi_2$  must be satisfied.

**Flat-9** Given  $\varphi = F_{[a;b]}(\varphi_1)$ , for each constraint  $(l_1; h_1; p_1) \geq X_1^\circ$ , flattening involves shifting interval  $[l_1; h_1]$  by a symbolic time variable  $t \geq [a; b]$ . This

bound on  $t$  is encoded by adding a tuple  $(a; b; t)$  to  $TC$ . Intuitively, this can be understood from the following transformation similar to the one in [13]:

$$F_{[a;b]}(F_{[c;d]}(\varphi)) = F_{[a+c; b+d]}(\varphi)$$

Here, instead of adding  $a$  and  $b$  directly to the upper and lower bound, respectively of time interval  $[c; d]$ , we add a symbolic variable that can take a value between  $a$  and  $b$ . For example, given specification  $\varphi = F_{[2;5]}(p)$ ,  $X^\varphi$  would contain constraint  $(t+0; t+0; p)$  and  $TC$  would contain  $(2; 5; t)$ .

As another example, consider specification  $\varphi = F_{[2;5]}(p_2 \wedge F_{[1;3]}(p_1))$ . When the  $\text{Flat-}\varrho$  rule is applied, the resulting  $X^\varphi$  contains  $(t_2 + t_1 + 0; t_2 + t_1 + 0; p_1)$  and  $(t_2; t_2; p_2)$ , and  $TC$  contains  $f(1; 3; t_1)$ ,  $(2; 5; t_2)g$ . Here,  $t_1$  is introduced when  $\text{Flat-}\varrho$  rule is applied to the innermost  $F$  operator ( $F_{[1;3]}$ );  $t_2$  is then introduced when  $\text{Flat-}\varrho$  is applied for outermost  $F$  operator ( $F_{[2;5]}$ ).

The flattening of an  $F$  formula applied to an invariance set is handled similarly. Consider a specification of the form

$$F_{[a;b]}(\varphi) = F_{[a;b]}(G_{[c;d]}(p))$$

with  $X_1^\varphi = f(c; d; p)g$ . Here, according to the Boolean STL semantics, there exists a  $t \geq [a; b]$  such that  $\delta t^\theta \geq [t + c; t + d]$ ,  $p$  must be satisfied. Hence, in the resulting invariance set for  $\varphi$ , the time interval for the existing invariance constraint is shifted by  $t$ . This bound on  $t$  is encoded by adding a tuple  $(a; b; t)$  to  $TC$ . For example, given specification  $\varphi = F_{[2;5]}(G_{[3;10]}(p))$ ,  $X^\varphi$  would contain constraint  $(t + 3; t + 10; p)$  and  $TC$  would contain  $(2; 5; t)$ .

**Flat- $\delta$**  The flattening of an  $\delta$  formula over  $X_1^\varphi$  is handled in the following way. Consider a specification of the form:

$$G_{[a;b]}(\varphi) = G_{[a;b]}(F_{[c;d]}(p))$$

with  $X_1^\varphi = ftcg = f(t_1; t_1; p)g$  and  $TC = f(c; d; t_1)g$ . Since  $\varphi$  states that constraint  $tc$  must hold at every time step between  $[a; b]$ , the idea is to create multiple reachability constraints of form  $(t_1 + k; t_1 + k; p)$ , one for each time value  $k$  in interval  $[a; b]$ . For example, let  $\varphi = G_{[1;100]}(F_{[1;5]}(p))$ . After flattening,  $X^\varphi$  would contain 100 reachability constraints, each in form of  $(t_1 + k; t_1 + k; p)$  for  $1 \leq k \leq 100$ .

The flattening of an  $\delta$  over  $X_1^\varphi$  is handled in the following way. Consider a specification of the form

$$G_{[a;b]}(\varphi) = G_{[a;b]}(G_{[c;d]}(p))$$

with  $X_1^\varphi = ftcg = f(c; d; p)g$ . Since  $\varphi$  states that constraint  $tc$  must hold at every time step between  $[a; b]$ , by shifting the interval of each constraint by  $[a; b]$ , the resulting formula  $X^\varphi$  is constructed as  $f(a + c; b + d; p)g$ . Intuitively, this can be understood from the following transformation similar to the one in [13]:

$$G_{[a;b]}(G_{[c;d]}(\varphi)) = G_{[a+c; b+d]}(\varphi)$$



We now introduce a theorem that establishes the soundness of our flattening operation with respect to a given STL specification.

**Lemma 1.** *Let  $\phi$  be an input STL specification, and let  $(X^{\mathcal{R}}, X^{\mathcal{I}}, TC)$  be the output of  $\text{flatten}(\phi)$ . Then, for every signal  $s_t$ :*

$$\exists v \geq V(\text{vars}(TC)) \quad \exists x \geq X^{\mathcal{R}} [X^{\mathcal{I}} \quad s_t \models \text{Inst}(x; v) \Rightarrow \quad s_t \models \phi]$$

where  $\text{vars}$  is a function that returns the set of all symbolic variables in  $TC$ ,  $V(X)$  is the set of all possible assignments of values to variables in  $X$  (restricted to values in their respective time intervals), and  $\text{Inst}(x; v)$  instantiates symbolic variables in constraint  $x$  with the values from  $v$ .

In other words, if every possible reachability or invariance constraint is satisfied over a particular signal, then the original specification  $\phi$  must also hold over the same signal.

### 4.3 Symbolic Time Resolution

The  $X^{\mathcal{R}}$  and  $X^{\mathcal{I}}$  constraints generated from flattening may contain multiple symbolic variables. To enable scheduling, `StlInc` first attempts to resolve as many of these variables as possible by applying substitution rules to the constraints.

Algorithm 1 shows the sketch of the *symbolic time resolution (STR)* process. As inputs, it takes the output of the flattening procedure—the  $X^{\mathcal{R}}$  and  $X^{\mathcal{I}}$  constraints, and the time intervals over symbolic time variables. As outputs, it produces (1) a new pair of  $X^{\mathcal{R}}$  and  $X^{\mathcal{I}}$ , with some of the time variables replaced by concrete time values, and (2) time variable intervals ( $TC^{\mathcal{I}}$ ) that specifies, for each concrete constraint produced in (1), an interval during which the constraint must be satisfied. This latter set of time intervals are used to enforce conjunctive constraints (i.e.,  $\phi_1 \wedge \phi_2$ ) to be satisfied simultaneously.

The algorithm iterates through the time variables in the bottom-up order (i.e., starting with the variables that appear in the lower part of an AST for a given STL expression). For each constraint that contains variable  $t$  (line 4), STR applies two types of substitutions, depending on whether the constraint is a reachability or an invariance constraint.

**Invariance constraints** An invariance constraint,  $x = (l; h; \rho)$ , with time variable  $t \geq [a; b]$  represents an STL expression  $F_{[a; b]}(G_{[l; h]} \rho)$  (where  $t$  appears in  $l$  and  $h$ , and  $\rho$  is an atomic proposition). Intuitively, this expression can be regarded a kind of “reach and stay” task, where the system must first reach a state where  $\rho$  holds within time interval  $[a + l; b + l]$ , and then continue to satisfy  $\rho$  for the following  $(h - l)$  time steps. The rule `ApplyFG` takes constraint  $x$  and interval  $tc$ , and produces a new pair of reachability and invariance constraints, as follows:

$$\text{ApplyFG}(x = (l; h; \rho); tc = (a; b; t)) \quad ((a + l; b + l; \rho); (t_{\text{sat}}; t_{\text{sat}} + h - l; \rho))$$

**Algorithm 1** Symbolic Time Resolution (STR)

---

```

1: Input: Flattened constraints  $X^\circ; X^\beta$ , time variable intervals  $TC$ 
2: Output: Modified constraints  $X^\circ; X^\beta$ , new time variable intervals  $TC^\circ$ 
3:  $TC^\circ = ;$ 
4: for  $tc = (a; b; t)$  in  $TC$  do
5:    $X_t^\beta := \text{filter}(X^\beta; t)$ ,  $X_t^\circ := \text{filter}(X^\circ; t)$ 
6:   for  $x = (l; h; \rho)$  in  $X_t^\beta$  do
7:      $(y^\circ; y^\beta) := \text{ApplyFG}(x, tc)$ 
8:      $X^\circ := X^\circ [f y^\circ g$ 
9:      $X^\beta := (X^\beta \quad f x g) [f y^\beta g$ 
10:     $TC^\circ := TC^\circ [f(t+l; t+h; \text{t-sat}(y^\circ)); (t+l+1; t+h; \text{t-sat}(y^\beta))g$ 
11:   end for
12:   for  $x = (l; h; \rho)$  in  $X_t^\circ$  do
13:      $y^\circ := \text{ApplyFF}(x, tc)$ 
14:      $X^\circ := (X^\circ \quad f x g) [f y^\circ g$ 
15:      $TC^\circ := TC^\circ [f(t+l; t+h; \text{t-sat}(y^\circ))g$ 
16:   end for
17: end for
18: return  $(X^\circ; X^\beta; TC^\circ)$ 

```

---

where  $t_{sat}$  is a new symbolic variable that represents the time at which  $\rho$  is satisfied between  $(a+l; b+l)$ .


In addition, STR adds two time variable intervals (line 10) to ensure that: (1) the new reachability constraint,  $y^\circ$ , is satisfied exactly at  $(t+l)$  and (2) the invariance constraint,  $y^\beta$ , is satisfied subsequently for the following  $(l-h)$  steps. Here,  $\text{t-sat}(y)$  returns a symbolic time variable representing the time of satisfaction of constraint  $y$ .

Consider the example in Figure 4. One of the invariance constraints that flattening generates is  $(t+1; t+5; r_1)$ , which depends on time variable  $t \in [1; 20]$ . The application of `ApplyFG` results in constraints stating that (1) the system must satisfy  $\rho$  within  $[2; 21]$ , and (2) from the point of the satisfaction of this constraint ( $t_2$ ), it must hold  $\rho$  true for the following  $(5-1) = 4$  steps. Note that the other invariance constraint,  $(1; 35; r_3)$ , remains untouched, as it does not depend on any time variable.

**Reachability constraints** A reachability constraint,  $x = (l; h; \rho)$ , with time variable  $t \in [a; b]$  represents an STL expression  $F_{[a;b]}(F_{[l;h]} \rho)$ . When a pair of  $F$  operators are nested in this manner, they can be simplified by using the following substitution rule:

$$\text{ApplyFF}(x = (l; h; \rho); tc = (a; b; t)) \quad (l+a; h+b; \rho)$$

The resulting reachability constraint,  $y^\circ$ , replaces the existing constraint in  $X^\circ$  (line 14). In addition, a new time variable interval is added to ensure that  $y^\circ$  is satisfied between  $(t+l)$  and  $(t+h)$ .



**Fig. 4.** An AST for an example STL specification,  $F_{[1,20]}(G_{[1,5]}(r_1) \wedge F_{[6,15]}(r_2)) \wedge G_{[1,35]}(\neg r_3)$ , with outputs from the flattening and symbolic time resolution steps. Each constraint resulting from the resolution step is assigned an identifier ( $id$ ); a symbolic time variable that represents the time of satisfying the constraint is annotated with  $id$  as the subscript (i.e.,  $t_2$  represents the satisfaction of reachability constraint  $(2;21;r_1)$ ).

For the example in Figure 4, flattening produces one reachability constraint,  $(t+t^\ell; t+t^\ell; r_2)$ . During the first iteration of the outermost loop (line 4), STR selects  $tc = (5;15;t^\ell)$  and applies `ApplyFF` to the constraint, producing a new reachability constraint  $(t+5; t+15; r_2)$ . In next iteration, STR selects  $tc = (1;20;t)$  and applies `ApplyFF` to  $(t+5; t+15; r_2)$ , producing an additional reachability constraint  $(6;35;r_2)$ .

Note that at the end of resolution for the example (Figure 4), variable  $t$  appears in multiple intervals in  $TC^\theta$ . In the following section, we describe how `StlInc` schedules tasks to generate concrete values for the time variables incrementally one-by-one, ultimately synthesizing a plan that satisfies the original STL formula.

**Lemma 2.** Let  $(X^\theta, X^\beta, TC)$  be the output of `flatten( )` and  $(X_2^\theta, X_2^\beta, TC^\theta)$  be the output of `STR( $X^\theta; X^\beta; TC$ )`. Then, for every signal  $s_t$ :

$$\begin{aligned} \exists v \geq \forall (\text{vars}(TC^\theta)) \quad & (\exists x_2 \geq X_2^\theta [ X_2^\beta \quad s_t \neq \text{Inst}(x_2; v) ] \Rightarrow) \\ & (\exists x \geq X^\theta [ X^\beta \quad s_t \neq \text{Inst}(x; v) ]) \end{aligned}$$

In other words, if every reachability and invariance constraint generated from STR is satisfied under some instantiation ( $v$ ) of symbolic time variables in  $TC^\theta$ , then the original set of flattened constraints must also be satisfied under the same condition.

#### 4.4 Scheduling

Given the output from the resolution step (the two constraint sets,  $X^\theta; X^\beta$  and the time interval variables,  $TC^\theta$ ), the goal of the scheduler is synthesize a plan that satisfies the original STL specification . To achieve this, the scheduler iteratively interacts with a *planner* that is capable of synthesizing a plan to

**Algorithm 2** Schedule

---

```

1: Input:  $X^{\mathcal{Q}}; X^{\mathcal{B}}; TC^0$ 
2: Output: Signal  $s_{plan}$  for synthesized plan
3:  $s_{plan} = hi$ 
4:  $\prec := \text{ComputeOrder}(X^{\mathcal{Q}}; X^{\mathcal{B}}; TC^0)$ 
5:  $currTasks := \text{NextTasks}(X^{\mathcal{Q}}; X^{\mathcal{B}}; \prec; ;)$ 
6: while  $currTasks \neq \emptyset$  do
7:    $currTasks := \text{Slice}(currTasks)$ 
8:    $atomicTask := \text{NextAtomic}(currTasks)$ 
9:    $currTasks := currTasks \setminus atomicTask$ 
10:   $s_{atomic} := \text{Plan}(atomicTask)$ 
11:  if  $s_{atomic} = hi$  then break end if
12:   $s_{plan} := s_{plan} \cup s_{atomic}$ 
13:   $satTasks := \text{ExtractTime}(s_{atomic})$ 
14:   $currTasks := currTasks \cup \text{NextTasks}(X^{\mathcal{Q}}; X^{\mathcal{B}}; \prec; satTasks)$ 
15: end while
16: return  $s_{plan}$ 

```

---

satisfy an *atomic task* formula of form  $F_{[a;b]}(\varphi_1) \wedge G_{[a;b]}(\varphi_2)$ , where  $\varphi_1$  and  $\varphi_2$  are quantifier-free STL expressions. The scheduling process is incremental: The scheduler generates a sequence of atomic tasks formulas and invokes the planner to solve them one-by-one, using information (i.e., the time of satisfaction of a reachability or invariance constraint) generated by the planner to resolve any dependencies on symbolic time variables that were introduced during the flattening and resolution steps. The scheduling algorithm (Alg. 2) comprises of three major parts: ordering, slicing of constraints, and planning of atomic tasks.

**Ordering** In the first step (line 4), the scheduler computes a partial order  $\prec$  among the given reachability and invariance constraints, to determine which of these constraints must be satisfied before others. In particular, given a pair of constraints,  $\chi_1$  and  $\chi_2$ ,  $\chi_1 \prec \chi_2$  if and only if the time of satisfaction of  $\chi_1$  necessarily precede that of  $\chi_2$ , based on the time intervals that are assigned to those constraints in  $X^{\mathcal{Q}}$  or  $X^{\mathcal{B}}$ .

If one or more of  $\chi_1$  and  $\chi_2$  depends on another symbolic variable,  $t$ , for satisfaction, then the information in  $TC^0$  is used to determine the presence of a precedence relationship. Consider  $t_1$  and  $t_3$  from Figure 4; after resolution, the satisfaction of these two constraints depend on the symbolic variable  $t$  (as specified in  $TC^0$ ). Although the value of  $t$  is unknown, it can be determined that for any possible value of  $t$ ,  $t_3$  will necessarily be satisfied before  $t_1$  (i.e.,  $t_3 \prec t_1$ ). Overall, for this example, `ComputeOrder` determines that  $t_2 \prec t_3 \prec t_1$ . Note that  $t_4$  does not appear in this ordering as it needs to be satisfied in parallel with these other constraints.

**Slicing** The scheduler then determines the first set of constraints (or tasks) to be carried out based on the order  $\prec$  (line 5). In general, the time intervals over these constraint may overlap with each other in an arbitrary way. Recall,

however, that each atomic task to be solved by the planner must be in form  $F_{[a;b]}(\tau_1) \wedge G_{[a;b]}(\tau_2)$ . Thus, the scheduler must first convert the constraints in  $currTasks$  into a set of atomic tasks constraint; this step involves *slicing* one or more constraints in  $currTasks$  (line 7).

Due to limited space, we provide the full details of the slicing algorithm in the appendix. We briefly illustrate it here using an example, shown in Figure 5. For our running example, the first set of constraints to be fulfilled is  $f\tau_2 = (2;21;r_1); t_4 = (1;35;r_3)g$ . To generate atomic tasks out of these, the Slice operation slices  $t_4$  into three constraints,  $t_{4a}; t_{4b}; t_{4c}$ ; this, in turn, results in the following three atomic tasks:




Fig. 5. A slicing example.

$$G_{[1;1]}(\tau_3) \quad F_{[2;21]}(\tau_2) \wedge G_{[2;21]}(\tau_3) \quad G_{[22;35]}(\tau_3)$$

However, the slicing step may introduce dependencies among the atomic tasks, especially for the constraints in  $X^g$ . For example, suppose that task  $F_{[2;21]}(\tau_2)$  is further split into two slices,  $F_{[2;12]}(\tau_2)$  and  $F_{[12;21]}(\tau_2)$ . Since  $\tau_2$  needs to be satisfied only once in interval  $[2;21]$ , we would need to ensure that we are not over-constraining the space of possible behaviors by requiring  $\tau_2$  to be satisfied twice, as the slices would imply. To achieve this, we keep track of dependencies among constraints, which are used by the scheduler to remove unnecessary atomic tasks (e.g., remove  $F_{[12;21]}(\tau_2)$  once  $\tau_2$  is satisfied between  $[2;12]$ ). This dependency management is handled inside NextTasks (Algorithm 2, Line 14).

**Planning atomic tasks** Once atomic tasks have been generated through slicing, the scheduler selects the next atomic task and invokes the planner (lines 8-10). If the planner is able to synthesize a plan that satisfies the atomic task, it returns a signal that represents the satisfying trajectory, which is then appended to the cumulative signal  $S_{plan}$  (line 12); if not, the scheduler terminates by returning the signal that contains a partially satisfactory plan (line 11).

From the synthesized signal, the scheduler extracts the constraints from  $X^g$  and  $X^b$  that were satisfied, along with the concrete time values for their satisfaction (line 13). This information ( $satTasks$ ) is then used to determine the next increment of constraints to be solved.

Finally, when all of the constraints in  $X^g$  and  $X^b$  have been satisfied, the scheduler terminates by returning  $S_{plan}$  as the final plan.

**Lemma 3.** *Let  $S_{plan}$  be the output of  $Schedule(X^g, X^b, TC^0)$  and  $v$  be an assignment of values to symbolic time variables in  $TC^0$ , which are determined during the scheduling step. Then, the following statement holds:*

$$\exists x \exists X^g [ X^b \ S_{plan} \models Inst(x; v)$$

In other words, the synthesized plan,  $S_{plan}$ , satisfies all of the reachability and invariance constraints that were generated from the preceding flattening and STR steps.

	STL Specification	Pattern
1	$F_{[0:15]}(R_1) \wedge F_{[5:25]}(R_2) \wedge F_{[20:30]}(R_3) \wedge G_{[0:40]}(: O_1)$	R+A
2	$F_{[0:15]}(R_1 \wedge F_{[0:15]}(R_2)) \wedge G_{[0:40]}(: O_1)$	SV+A
3	$F_{[0:15]}(R_1 \wedge F_{[0:15]}(R_2 \wedge F_{[0:20]}(R_3 \wedge F_{[0:15]}(R_1))))$	SV
4	$F_{[0:15]}G_{[0:10]}(R_1) \wedge F_{[0:35]}(R_2) \wedge G_{[0:40]}(: O_1)$	R+A+SB
5	$F_{[0:15]}(R_1 \wedge F_{[0:20]}G_{[0:10]}(R_2))$	SV+SB

**Table 1.** Benchmark STL specifications created from motion planning patterns. Here, R: Reach, A: Avoid, SV: Sequenced Visit, SB: Stabilization.

Building on Lemmas 1, 2, and 3, we finally introduce a theorem to state that our proposed approach generates a plan that satisfies the given specification :

**Theorem 1.** *Given specification  $\phi$  and  $S_{plan}$  as the output of the scheduling algorithm,  $S_{plan} \models \phi$ .*

## 5 Evaluation

This section begins with a detailed description of our experimental setup, including specifications and implementation details, alongside the research questions we aim to address. Following this, we present our findings and conclude with a discussion of our approach over the benchmarks.

### 5.1 Experimental Setup

**Specifications.** We investigated multiple motion planning STL specifications from [6, 17]. Based on the most common planning patterns, such as Reach (R), Avoid (A), Stabilisation (SB), Sequenced Visits (SV) etc., we created representative STL benchmark specifications as outlined in Table 1. These specifications are defined over STL subformulas of the form  $R_i$  or  $O_i$  where  $R_i / O_i$  is satisfied if the agent is inside Region  $i$  or Obstacle  $i$ . These subformulas are defined in a similar fashion using conjunction of linear and nonlinear predicates as done in [12]. Please refer to [12] for more information on how these are defined for rectangular/circular regions.

**Implementation Details.** We investigate planning from benchmark specifications in two robot exploration environments (similar to Figure 1), namely LinEnv and NonLinEnv created using `stlpy` [12]. `stlpy` has the functionality to encode any arbitrary STL formula, dynamics and actuation limits into constraints and use existing state-of-the-art solvers (Gurobi [10], SNOPT [9], etc.) to generate satisfying plans. Our two environments are both planar but differ in underlying dynamics governing the robot. LinEnv has linear dynamics (Double Integrator) whereas NonLinEnv has nonlinear dynamics (Unicycle). We benchmark our technique against existing MICP methods (for linear dynamics) and

other gradient-based techniques like SNOPT (for non-linear dynamics). We use Python to implement our tool<sup>4</sup> while using `stlpy` and Drake [23] to encode the STL constraints. Additionally, we use Gurobi or SNOPT to solve the final constrained optimization problem. All experiments were run on a workstation with an Intel Xeon W-1350 processor and 32 GB RAM.

**Benchmarks and Research Questions.** We compare against the state-of-the-art techniques proposed in [4] (which we call *standard MICP*) and [12] (*reduced MICP*). Since the standard MICP encoding is only defined for environments with linear dynamics, we compare our technique against reduced MICP encoding for NonLinEnv. Reduced MICP claims better performance over standard MICP for long horizon and complex specifications due to their efficient encoding of disjunction and conjunction with fewer binary variables. However, standard MICP is faster for short-horizon specifications due to solver-specific presolve routines that leverage the additional binary variables for simplification.

Since our focus is on both short- and long-horizon specification with deep levels of temporal operator nesting, we benchmark against both techniques. The two main metrics we are concerned with are the time taken for solving and the final robustness values. To make the comparison fair, the total time taken by our technique includes the time taken by the flattener, scheduler, and solvers.

The two main research questions we investigate in this paper are:

1. **RQ1:** Does our decomposition technique result in shorter solve times?
2. **RQ2:** Does our decomposition technique result in higher robustness scores?

## 5.2 Results

Table 2 summarizes the results for `StlInc` performance compared to the baselines. In the tables, `N` represents the horizon of the specification and `D` represents the maximum depth of temporal nesting; `TO` represents a timeout, which means the solver did not terminate despite running it for 30 minutes. In those cases, the solver’s output plan robustness is represented as `-inf` (which means no solution was found in the given time).

For `LinEnv` for all the specifications, our robustness values are comparable to the two techniques but our solve times are either lower or comparable to the baselines. Additionally, for specification `s3`, which has the deepest temporal nesting, our method significantly outperforms both baseline methods that experience timeouts.

For `NonLinEnv` for `s1` and `s4`, the baseline encoding performs better in terms of solving time but `StlInc` only does slightly worse. However, for specification `s2`, `s3` and `s5`, `StlInc` significantly outperforms the baselines.

<sup>4</sup> <https://github.com/parvkpr/MCTSTL>

Spec	N	D	Solve time (s)					Robustness				
			LinEnv			NonLinEnv		LinEnv			NonLinEnv	
			[4]	[12]	StlInc	[12]	StlInc	[4]	[12]	StlInc	[12]	StlInc
1	40	0	<b>0.845</b>	2.698	0.891	<b>0.890</b>	1.464	0.500	0.500	<b>0.500</b>	<b>0.430</b>	0.572
2	30	1	2.459	TO	<b>0.402</b>	12.674	<b>0.892</b>	0.491	-inf	<b>0.491</b>	-inf	<b>0.594</b>
3	60	3	TO	TO	<b>0.874</b>	15.829	<b>1.554</b>	-inf	-inf	<b>0.228</b>	-inf	<b>0.065</b>
4	40	2	<b>0.318</b>	0.330	0.629	<b>1.049</b>	1.131	<b>0.494</b>	0.500	0.500	0.470	<b>0.364</b>
5	40	2	2.829	28.490	<b>0.694</b>	83.193	<b>1.776</b>	0.500	0.500	<b>0.500</b>	-inf	<b>0.596</b>

**Table 2.** StlInc Performance Benchmarking for LinEnv and NonLinEnv against standard MICP ([4]) and reduced MICP ([12]).

**Summary.** Our technique excels significantly for nesting depths  $> 1$  in both LinEnv and NonLinEnv. However, for nesting depths  $< 1$ , the baseline techniques outperform us due to marginal overhead from flattening, scheduling, and solver invocations. The encoding for these specifications involves fewer binary variables, and the preprocessing overhead of using StlInc outweighs the performance benefits. Nevertheless, the experiment suggests that our technique is more efficient for multi-step tasks with deep temporal nesting, outperforming baselines by an order of magnitude.

## 6 Related Work

Trajectory synthesis from STL specifications is an active area of research for which multiple approaches have been proposed in the past few years [20, 2, 21, 1, 15]. One of the first papers in this direction involved translating STL specifications into constraints within a Mixed Integer Linear Program (MILP) [20]. This approach is sound and complete but faces scalability challenges for long-horizon specifications. To remedy this drawback, the original encoding has been modified by focusing on abstraction-based techniques [22] and reducing binary variables via logarithmic encoding [12]. Most of these techniques focus on reducing the MILP’s complexity to observe performance benefits. Recently, the focus has shifted to developing techniques that leverage robustness feedback as a heuristic for trajectory synthesis instead of using MILP. These techniques involve using reinforcement learning [1, 11], search-based techniques [3] and control barrier functions [15] to generate STL satisfying trajectories. While these methods offer greater scalability, they are not complete and frequently struggle to accommodate complex specifications because of the intrinsically non-convex optimization problem posed by robustness semantics.

In this work, we focus on MILP-based techniques due to their completeness guarantees and improve their scalability by modifying input STL specifications themselves. However, since we perform structural manipulation of the specifications themselves, our decomposition technique is planner-agnostic and can also use learning- or search-based planners. Decomposition of STL specifications has



been studied before in [13, 24]. However, our work differs from existing work in multiple ways. In [24], the authors restrict themselves to an STL fragment that does not allow nesting of temporal operators, while a key contribution of our work is handling deep nested specifications. In [13], the authors perform structural manipulation using a tree structure. However, their focus is on multi-agent setups and they handle nested operators conservatively, especially for the eventually operator. This conservative notion generates specification satisfying behavior but it can be overly restrictive. Our interpretation is more flexible and in line with the Boolean semantics defined for the same operators.

## 7 Limitations and Future Work

In this work, we propose a structural manipulation-based technique for the temporal decomposition of STL specifications, enabling the incremental fulfillment of these specifications. We show our method generates correct-by-construction trajectories that satisfy deeply nested specifications with long time horizons for which existing baseline STL planning techniques struggle.

While the proposed approach is promising, our current decomposition technique does not handle disjunction. Additionally, the technique is sound but not complete, and designed to prioritize satisfaction over optimality. This limitation stems from incremental planning of objectives, which can be locally optimal compared to global planning, which considers the entire problem space. In future work, we aim to enrich our scheduling algorithm with backtracking capability, which can generate multiple satisfying plans, to overcome this limitation. Furthermore, for probabilistic systems, we plan to employ conformal prediction techniques to overcome compounding modeling error issues. Another avenue for future work is to adapt our theory to accommodate learning and heuristic-based approaches, such as reinforcement learning. Finally, our decomposition theory can potentially be employed for other STL applications, such as falsification, testing and runtime assurance and we plan to investigate that in the future.

## Acknowledgments

We'd like to thank our reviewers for their insightful feedback. This work was supported in part by the National Science Foundation Award CCF-2144860.

## References

1. Aksaray, D., Jones, A., Kong, Z., Schwager, M., Belta, C.: Q-learning for robust satisfaction of signal temporal logic specifications (2016)
2. Aloor, J.J., Patrikar, J., Kapoor, P., Oh, J., Scherer, S.: Follow the rules: Online signal temporal logic tree search for guided imitation learning in stochastic domains. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 1320–1326 (2023). <https://doi.org/10.1109/ICRA48891.2023.10160953>

3. Aloor, J.J., Patrikar, J., Kapoor, P., Oh, J., Scherer, S.: Follow the rules: On-line signal temporal logic tree search for guided imitation learning in stochastic domains. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 1320–1326. IEEE (2023)
4. Belta, C., Sadraddini, S.: Formal methods for control synthesis: An optimization perspective. *Annual Review of Control, Robotics, and Autonomous Systems* **2**(1), 115–140 (2019). <https://doi.org/10.1146/annurev-control-053018-023717>, <https://doi.org/10.1146/annurev-control-053018-023717>
5. Botea, A., Ciré, A.A.: Incremental heuristic search for planning with temporally extended goals and uncontrollable events. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence. p. 1647–1652. IJCAI’09, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2009)
6. Chen, Y., Gandhi, R., Zhang, Y., Fan, C.: NL2TL: transforming natural languages to temporal logics using large language models. *CoRR* **abs/2305.07766** (2023). <https://doi.org/10.48550/arXiv.2305.07766>, <https://doi.org/10.48550/arXiv.2305.07766>
7. Czechowski, K., Odrzygóźdź, T., Zbysiński, M., Zawalski, M., Olejnik, K., Wu, Y., Kuciński, L., Miłoś, P.: Subgoal search for complex reasoning tasks. *Advances in Neural Information Processing Systems* **34**, 624–638 (2021)
8. Donnarumma, F., Maisto, D., Pezzulo, G.: Problem solving as probabilistic inference with subgoaling: Explaining human successes and pitfalls in the tower of hanoi. *PLoS Computational Biology* **12**(4) (2016)
9. Gill, P.E., Murray, W., Saunders, M.A.: Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM Rev.* **47**(1), 99–131 (jan 2005). <https://doi.org/10.1137/S0036144504446096>, <https://doi.org/10.1137/S0036144504446096>
10. Gurobi Optimization, Inc.: Gurobi optimizer reference manual (2012), <http://www.gurobi.com>
11. Kapoor, P., Balakrishnan, A., Deshmukh, J.V.: Model-based reinforcement learning from signal temporal logic specifications. *arXiv preprint arXiv:2011.04950* (2020)
12. Kurtz, V., Lin, H.: Mixed-integer programming for signal temporal logic with fewer binary variables. *IEEE Control Systems Letters* (2022)
13. Leahy, K., Mann, M., Vasile, C.I.: Rewrite-based decomposition of signal temporal logic specifications. In: Rozier, K.Y., Chaudhuri, S. (eds.) *NASA Formal Methods*. pp. 224–240. Springer Nature Switzerland, Cham (2023)
14. Leung, K., Aréchiga, N., Pavone, M.: Backpropagation through signal temporal logic specifications: Infusing logical structure into gradient-based methods. *The International Journal of Robotics Research* p. 02783649221082115 (2023)
15. Lindemann, L., Dimarogonas, D.V.: Control barrier functions for signal temporal logic tasks. *IEEE Control Systems Letters* **3**, 96–101 (2019), <https://api.semanticscholar.org/CorpusID:50767137>
16. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: *FORMATS/FTRIFT* (2004), <https://api.semanticscholar.org/CorpusID:15642684>
17. Menghi, C., Tsigkanos, C., Pelliccione, P., Ghezzi, C., Berger, T.: Specification patterns for robotic missions. *IEEE Transactions on Software Engineering* **47**(10), 2208–2224 (2021). <https://doi.org/10.1109/TSE.2019.2945329>
18. Nair, S., Finn, C.: Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. *ArXiv* **abs/1909.05829** (2019), <https://api.semanticscholar.org/CorpusID:202565422>

19. Pnueli, A.: The temporal logic of programs. In: 18th Annual Symposium on Foundations of Computer Science (sfcs 1977). pp. 46–57 (1977). <https://doi.org/10.1109/SFCS.1977.32>
20. Raman, V., Donzé, A., Maasoumy, M., Murray, R.M., Sangiovanni-Vincentelli, A., Seshia, S.A.: Model predictive control with signal temporal logic specifications. In: 53rd IEEE Conference on Decision and Control. pp. 81–87 (2014). <https://doi.org/10.1109/CDC.2014.7039363>
21. Sadraddini, S., Belta, C.: Formal synthesis of control strategies for positive monotone systems. *IEEE Transactions on Automatic Control* **64**(2), 480–495 (2019). <https://doi.org/10.1109/TAC.2018.2814631>
22. Sun, D., Chen, J., Mitra, S., Fan, C.: Multi-agent motion planning from signal temporal logic specifications. *IEEE Robotics and Automation Letters* **PP**, 1–1 (2022), <https://api.semanticscholar.org/CorpusID:245986629>
23. Tedrake, R., the Drake Development Team: Drake: Model-based design and verification for robotics (2019), <https://drake.mit.edu>
24. Yu, X., Wang, C., Yuan, D., Li, S., Yin, X.: Model predictive control for signal temporal logic specifications with time interval decomposition (2022)

## A Appendix

### A.1 Slicing Algorithm and Illustration

Our slicing algorithm is illustrated in Algorithm 3 and its application is demonstrated using the running example. In this example, the first set of constraints to be fulfilled is  $f t_2 = (2;21;r_1); t_4 = (1;35;r_3)g$ . In the first step (line 3), we identify the lowest and the highest time bounds out of all the constraints in *currTasks* (which, for the example, would be 1 and 35). Then, in lines 5 to 11, for the horizon, we identify which constraints are active at each given time step. After this step, in lines 14 to 29, we first create “slices”, which involves generating multiple time intervals out of time steps where similar constraints are active. Then, we create constraints of type  $X^s$  and  $X^g$  out of them by combining propositions of constraints of the same type. For the running example, these time slices would be [1,1], [2,21] and [22,35]. Finally, the constraints are converted into the following three atomic tasks:

$$G_{[1;1]}(: r_3) \quad F_{[2;21]}(r_2) \wedge G_{[2;21]}(: r_3) \quad G_{[22;35]}(: r_3)$$

---

**Algorithm 3** SLICE

---

```

1: Input: Set of currTasks  $X^c$ 
2: Output: Set of atomic tasks constraints  $X^{c^0}$ 
3:  $t_{min}, t_{max} := \text{Horizon}(\text{currTasks})$ 
4:  $active := hi$ 
5: for  $t$  in  $[t_{min}; t_{max}]$  do
6:    $active_t := hi$ 
7:   for  $x = (l; h; p)$  in  $X^c$  do
8:     if  $t$  in  $[l; h]$  then  $active_t := active_{\bar{t}} \wedge x$  end if
9:   end for
10:   $active := active \wedge active_t$ 
11: end for
12:  $t_{low} := t_{min}$ 
13:  $X^{c^0} = fg$ 
14: for  $t$  in  $[t_{min} + 1; t_{max}]$  do
15:   if  $active[t - 1] \neq active[t]$  then
16:      $X^{temp} := active[t - 1]$ 
17:      $X_{\varphi}^{temp} := f(t_{low}; t - 1; >)g$ 
18:      $X_{\beta}^{temp} := \bar{f}(t_{low}; t - 1; >)g$ 
19:     for  $x = (l; h; p)$  in  $X^{temp}$  do
20:       if  $x$  in  $X_{\varphi}$  then
21:          $X_{\varphi}^{temp}.p := X_{\varphi}^{temp}.p \wedge p$ 
22:       else
23:          $X_{\beta}^{temp}.p := X_{\beta}^{temp}.p \wedge p$ 
24:       end if
25:     end for
26:      $X^{c^0} := X^{c^0} [ X_{\varphi}^{temp} [ X_{\beta}^{temp}$ 
27:      $t_{low} := t$ 
28:   end if
29: end for
30: return  $X^{c^0}$ 

```

---