

Optimal Control of Nonlinear Systems with Temporal Logic Specifications

Eric M. Wolff and Richard M. Murray

Abstract We present a mathematical programming-based method for optimal control of nonlinear systems subject to temporal logic task specifications. We specify tasks using a fragment of linear temporal logic (LTL) that allows both finite- and infinite-horizon properties to be specified, including tasks such as surveillance, periodic motion, repeated assembly, and environmental monitoring. Our method directly encodes an LTL formula as mixed-integer linear constraints on the system variables, avoiding the computationally expensive process of creating a finite abstraction. Our approach is efficient; for common tasks our formulation uses significantly fewer binary variables than related approaches and gives the tightest possible convex relaxation. We apply our method on piecewise affine systems and certain classes of differentially flat systems. In numerical experiments, we solve temporal logic motion planning tasks for high-dimensional (10+ continuous state) systems.

1 Introduction

In safety-critical robotics applications involving autonomous ground and air vehicles, it is desirable to unambiguously specify the desired system behavior and automatically synthesize a controller that provably implements this behavior. Additionally, autonomous systems often have high-dimensional, nonlinear dynamics and require high-performance (not just feasible) controllers.

Linear temporal logic (LTL) is an expressive task-specification language for specifying a variety of tasks such as responding to the environment, visiting goals, periodically monitoring areas, staying safe, and remaining stable. These properties generalize classical point-to-point motion planning. Also, the widespread use of LTL

E.M. Wolff (✉) · R.M. Murray
California Institute of Technology, Pasadena, CA, USA
e-mail: ewolff@caltech.edu

R.M. Murray
e-mail: murray@cds.caltech.edu

in software verification [2] makes it appealing as a common language for reasoning about the software and dynamics of autonomous systems.

Standard methods for motion planning with LTL task specifications first create a finite abstraction of the original dynamical system. This abstraction can informally be viewed as a labeled graph that represents possible behaviors of the system. Approximate finite abstractions can be computed using either sampling-based methods (e.g., RRTs) [6, 14, 17] or reachability-based approaches [1, 3, 10, 16, 29].

Given a finite abstraction of a dynamical system and an LTL specification, controllers can be automatically constructed using an automata-based approach [2, 6, 9, 14, 16]. This approach first transforms the LTL formula into an equivalent Büchi automaton whose size may be exponential in the length of the formula [2]. A product automaton is created from the finite abstraction and the Büchi automaton, and then a controller is found by graph search in the product automaton.

The main drawback of this approach is that it is expensive to compute a finite abstraction. The product automaton might also be quite large due to the size of the abstraction and the Büchi automaton. Finally, although optimal controllers can be computed for the discrete abstraction [21, 27], optimality is only with respect to the abstraction’s level of refinement or asymptotic [14].

Instead of the automata-based approach, we directly encode a large class of temporal logic formulas as mixed-integer linear constraints on the original dynamical system. These constraints enforce that an infinite sequence of system states satisfies a task specification. A key component of our formulation is enforcing that the system is in a (non-convex) region at a given time. We introduce an alternative formulation for this that gives a tighter convex relaxation than the commonly used big-M approach. Our approach applies to any deterministic system model that is amenable to finite-dimensional optimization, as the temporal logic constraints are independent of any particular system dynamics or cost functions. We specifically investigate Mixed Logical Dynamic (MLD) systems [4] and certain differentially flat systems [19], whose dynamics can be encoded with mixed-integer linear constraints. MLD systems include constrained linear systems, linear hybrid automata, and piecewise affine systems. Differentially flat systems include quadrotors and car-like vehicles.

It is well-known that mixed-integer linear programming can be used for reasoning about propositional logic [7, 11], generating state-constrained trajectories [8, 20, 24], and modeling vehicle routing problems [13, 22]. The work most similar to ours is Karaman et al. [15], who consider controller synthesis for MLD systems subject to finite-horizon LTL specifications. However, finite-horizon properties are too restrictive to model a large class of interesting robotics problems, including persistent surveillance, repeated assembly, periodic motion, and environmental monitoring. Our work specifically addresses these types of periodic tasks with a novel mixed-integer formulation.

Our main contributions are (1) a novel method for encoding both finite- and infinite-horizon temporal logic properties as mixed-integer linear constraints on a system and (2) an improved encoding that has a tighter convex relaxation and uses significantly fewer binary variables for common tasks than related work [15]. The fragment of temporal logic that we consider allows one to specify properties

such as safety, stability, liveness, guarantee, and response. We demonstrate how this mixed-integer programming formulation can be used with off-the-shelf optimization solvers (e.g. CPLEX [23]) to compute both feasible and optimal controllers for high-dimensional systems with temporal logic specifications.

2 Preliminaries

An *atomic proposition* is a statement that is *True* or *False*. A *propositional formula* is composed of only atomic propositions and propositional connectives, i.e., \wedge (and), \vee (or), and \neg (not). Let $\mathcal{T} = \{0, 1, 2, \dots, T\} \subset \mathbb{N}$ denote a bounded set of discrete time instances and $\mathcal{T}^\infty = \{0, 1, 2, \dots\}$ denote an unbounded set of discrete time instances.

2.1 System Model

We consider discrete-time nonlinear systems of the form

$$x(t+1) = f(x(t), u(t)), \quad (1)$$

where $t \in \mathcal{T}^\infty$, $x \in \mathcal{X} \subseteq \mathbb{R}^{n_c} \times \{0, 1\}^{n_l}$ are the continuous and binary states, $u \in \mathcal{U} \subseteq \mathbb{R}^{m_c} \times \{0, 1\}^{m_l}$ are the inputs, and $x(0) = x_0 \in \mathcal{X}$ is the initial state. We assume that the system is deterministic, i.e., an initial state x_0 and a control input sequence $\mathbf{u} = u_0 u_1 u_2 \dots$ produces a unique trajectory (or run) $\mathbf{x} = \mathbf{x}(x_0, \mathbf{u}) = x_0 x_1 x_2 \dots$.

Let AP be a finite set of atomic propositions. The (time-dependent) *labeling function* $L_t : \mathcal{X} \rightarrow 2^{AP}$ maps the continuous part of each state to the set of atomic propositions that are *True* at time t . Each atomic proposition $\psi \in AP$ is represented by a union of polyhedrons. The finite index set I_t^ψ lists the polyhedrons where ψ holds at time t . The i th polyhedron is $\{x \in \mathcal{X} \mid H_t^{\psi_i} x \leq K_t^{\psi_i}\}$, where $i \in I_t^\psi$. Thus, the set of states where atomic proposition ψ holds at time t is given by $\llbracket \psi \rrbracket(t) := \{x \in \mathcal{X} \mid H_t^{\psi_i} x \leq K_t^{\psi_i} \text{ for some } i \in I_t^\psi\}$. This (potentially) time-varying set is the finite union of polyhedrons (finite conjunctions of halfspaces).

2.2 A Fragment of Temporal Logic

We do not attempt to reason about all possible temporal logic formulas (see [2]); instead, we develop a useful library of temporal operators for robotic tasks. This fragment of temporal logic can concisely and unambiguously specify a wide range of tasks such as safe navigation, surveillance, persistent coverage, response to the environment, and visiting goals. In the following definitions, ψ , ϕ , and ψ_j (for a

finite number of indices j) are propositional formulas. To simplify the presentation, we split these into three groups: core Φ_{core} , response Φ_{resp} , and fairness Φ_{fair} . We first define the syntax of the temporal operators and then their semantics.

Syntax

The core operators, $\Phi_{\text{core}} := \{\varphi_{\text{safe}}, \varphi_{\text{goal}}, \varphi_{\text{per}}, \varphi_{\text{live}}, \varphi_{\text{until}}\}$, specify fundamental properties such as safety, guarantee, persistence, liveness (recurrence), and until. These operators are,

$$\varphi_{\text{safe}} := \Box\psi, \quad \varphi_{\text{goal}} := \Diamond\psi, \quad \varphi_{\text{per}} := \Diamond\Box\psi, \quad \varphi_{\text{live}} := \Box\Diamond\psi, \quad \varphi_{\text{until}} := \psi \mathcal{U} \phi,$$

where φ_{safe} specifies safety, i.e., a property should invariantly hold, φ_{goal} specifies goal visitation, i.e., a property should eventually hold, φ_{per} specifies persistence, i.e., a property should eventually hold invariantly, and φ_{live} specifies liveness (recurrence), i.e., a property should hold repeatedly, as in surveillance, and φ_{until} specifies until, i.e., a property ψ should hold until another property ϕ holds.

The response operators, $\Phi_{\text{resp}} := \{\varphi_{\text{resp}}^1, \varphi_{\text{resp}}^2, \varphi_{\text{resp}}^3, \varphi_{\text{resp}}^4\}$, specify how the system responds to the environment. These operators are,

$$\begin{aligned} \varphi_{\text{resp}}^1 &:= \Box(\psi \implies \bigcirc\phi), & \varphi_{\text{resp}}^2 &:= \Box(\psi \implies \Diamond\phi), \\ \varphi_{\text{resp}}^3 &:= \Diamond\Box(\psi \implies \bigcirc\phi), & \varphi_{\text{resp}}^4 &:= \Diamond\Box(\psi \implies \Diamond\phi), \end{aligned}$$

where φ_{resp}^1 specifies next-step response to the environment, φ_{resp}^2 specifies eventual response to the environment, φ_{resp}^3 specifies steady-state next-step response to the environment, and φ_{resp}^4 specifies steady-state eventual response to the environment.

Finally, the fairness operators, $\Phi_{\text{fair}} := \{\varphi_{\text{fair}}^1, \varphi_{\text{fair}}^2, \varphi_{\text{fair}}^3\}$, allow one to specify conditional tasks. These operators are,

$$\begin{aligned} \varphi_{\text{fair}}^1 &:= \Diamond\psi \implies \bigwedge_{j=1}^m \Diamond\phi_j, & \varphi_{\text{fair}}^2 &:= \Diamond\psi \implies \bigwedge_{j=1}^m \Box\Diamond\phi_j, \\ \varphi_{\text{fair}}^3 &:= \Box\Diamond\psi \implies \bigwedge_{j=1}^m \Box\Diamond\phi_j, \end{aligned}$$

where φ_{fair}^1 specifies conditional goal visitation, and φ_{fair}^2 and φ_{fair}^3 specify conditional repeated goal visitation.

The fragment of LTL that we consider is built from the temporal operators defined above as follows,

$$\varphi := \varphi_{\text{core}} \mid \varphi_{\text{resp}} \mid \varphi_{\text{fair}} \mid \varphi_1 \wedge \varphi_2, \quad (2)$$

where $\varphi_{\text{core}} \in \Phi_{\text{core}}$, $\varphi_{\text{resp}} \in \Phi_{\text{resp}}$, and $\varphi_{\text{fair}} \in \Phi_{\text{fair}}$.

This LTL fragment specifies many properties relevant to robotics, especially for surveillance tasks for which no mathematical programming-based approaches currently exist. However, it does not include nested properties [2]. Determining all temporal properties that can be expressed in this framework is future work.

Remark 1 To include disjunctions (e.g., $\varphi_1 \vee \varphi_2$), one can rewrite a formula in disjunctive normal form, where each clause is of the form (2). In what follows, each clause can then be considered separately, as the system (1) is deterministic.

Semantics

We use set operations between a trajectory (run) $\mathbf{x} = \mathbf{x}(x_0, \mathbf{u})$ and subsets of \mathcal{X} where particular propositional formulas hold to define satisfaction of a temporal logic formula [2]. We denote the set of states where propositional formula ψ holds by $\llbracket \psi \rrbracket$. A run \mathbf{x} *satisfies* the temporal logic formula φ , denoted by $\mathbf{x} \models \varphi$, if and only if certain set operations hold. Given propositional formulas ψ and ϕ , we relate satisfaction of (a partial list of) formulas of the form (2) with set operations as follows:

- $\mathbf{x} \models \Box\psi$ iff $x_i \in \llbracket \psi \rrbracket$ for all i ,
- $\mathbf{x} \models \Diamond\Box\psi$ iff there exists an index j such that $x_i \in \llbracket \psi \rrbracket$ for all $i \geq j$,
- $\mathbf{x} \models \Diamond\psi$ iff $x_i \in \llbracket \psi \rrbracket$ for some i ,
- $\mathbf{x} \models \Box\Diamond\psi$ iff $x_i \in \llbracket \psi \rrbracket$ for infinitely many i ,
- $\mathbf{x} \models \psi \mathcal{U} \phi$ iff there exists an index j such that $x_j \in \llbracket \phi \rrbracket$ and $x_i \in \llbracket \psi \rrbracket$ for all $i < j$,
- $\mathbf{x} \models \Box(\psi \implies \bigcirc\phi)$ iff $x_i \notin \llbracket \psi \rrbracket$ or $x_{i+1} \in \llbracket \phi \rrbracket$ for all i ,
- $\mathbf{x} \models \Box(\psi \implies \Diamond\phi)$ iff $x_i \notin \llbracket \psi \rrbracket$ or $x_k \in \llbracket \phi \rrbracket$ for some $k \geq i$ for all i ,
- $\mathbf{x} \models \Diamond\Box(\psi \implies \bigcirc\phi)$ iff there exists an index j such that $x_i \notin \llbracket \psi \rrbracket$ or $x_{i+1} \in \llbracket \phi \rrbracket$ for all $i \geq j$,
- $\mathbf{x} \models \Diamond\Box(\psi \implies \Diamond\phi)$ iff there exists an index j such that $x_i \notin \llbracket \psi \rrbracket$ or $x_k \in \llbracket \phi \rrbracket$ for some $k \geq i$ for all $i \geq j$.

A run \mathbf{x} *satisfies* a conjunction of temporal logic formulas $\varphi = \bigwedge_{i=1}^m \varphi_i$ if and only if the set operations for each temporal logic formula φ_i holds. The LTL formula φ is *satisfiable* by a system at state $x_0 \in \mathcal{X}$ if and only if there exists a control input sequence \mathbf{u} such that $\mathbf{x}(x_0, \mathbf{u}) \models \varphi$.

3 Problem Statement

In this section, we formally state both a feasibility and an optimization problem and give an overview of our solution approach. Let φ be an LTL formula of the form (2) defined over AP .

Problem 1 Given a system of the form (1), with initial condition x_0 , and an LTL formula φ of the form (2), determine whether or not there exists a control input sequence \mathbf{u} such that $\mathbf{x}(x_0, \mathbf{u}) \models \varphi$.

We now introduce a cost function to distinguish among all trajectories that satisfy Problem 1. Since LTL formulas are defined over infinite state sequences, we define a cost function over infinite state sequences. We use a maximum cost function to simplify the presentation; it can easily be extended to discounted, limit-maximum, and average cost functions (see [26]). Let the cost $c : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ be bounded.

Definition 1 Let \mathbf{x} be a trajectory and \mathbf{u} be the corresponding control input sequence. The *maximum cost* of trajectory \mathbf{x} is

$$J(\mathbf{x}, \mathbf{u}) := \sup_{t \in T^\infty} c(x_t, u_t), \quad (3)$$

where J maps trajectories and control inputs to $\mathbb{R} \cup \infty$.

Problem 2 Given a system of the form (1), with initial condition x_0 , and an LTL formula φ of the form (2), compute a control input sequence \mathbf{u} such that $\mathbf{x}(x_0, \mathbf{u}) \models \varphi$ and $J(\mathbf{x}(x_0, \mathbf{u}), \mathbf{u})$ is minimized.

We now give a brief overview of our solution approach. We parameterize the system trajectory (control input) as a periodic prefix-suffix structure. Every LTL operator of the form (2) is encoded as mixed-integer linear constraints on this finite parameterization. These temporal logic constraints (see Sect. 5) are then combined with dynamic constraints (see Sect. 6) as constraints on a combined mixed-integer optimization problem with an appropriate cost function. For MLD systems and certain differentially flat systems (see Sect. 6) with linear costs, Problems 1 and 2 can thus be solved using a mixed-integer linear program (MILP) solver. While even checking feasibility of a MILP is NP-hard, modern solvers using branch and bound methods routinely solve large problems [23]. We show promising results (see Fig. 1) on high-dimensional (10+ continuous state) systems in Sect. 7.




Fig. 1 Illustration of a problem instance. The task is to repeatedly visit regions P , $D1$, and $D2$, where *dark* regions are obstacles that must be avoided. Representative trajectories for a quadrotor (*left*) and nonlinear car (*right*) are shown with the prefix (*blue*) and suffix (*black*)

Remark 2 We only consider open-loop trajectory generation, which is already a challenging problem due to the nonlinear dynamics and LTL specifications. Disturbances can be dealt with by wrapping a feedback controller around the trajectory. Incorporating disturbances during trajectory generation is the subject of future work.

4 A Periodic Trajectory Parameterization

We parameterize the system trajectory by a periodic prefix-suffix form that is commonly used in model checking for finite systems. In this structure, the prefix is a finite trajectory and the suffix is a finite trajectory that is repeated infinitely often. This gives a sufficient condition that is amenable to computation, although it may miss valid non-periodic trajectories.

A *walk* is a finite sequence of states $\mathbf{x} = x_0 x_1 x_2 \dots x_N$ that satisfy the constraints in (1). A *cycle* is a walk $\mathbf{x} = x_0 x_1 x_2 \dots x_N$ where $f(x_N, u) = x_0$ for some $u \in \mathcal{U}$. A trajectory \mathbf{x} induces a corresponding *word* (i.e., sequence of labels) $L(\mathbf{x}) = L_0(x_0)L_1(x_1)L_2(x_2) \dots$ through the labeling function. A word is similarly defined for a walk or cycle. We now define a trajectory in *prefix-suffix* form.

Definition 2 Let \mathbf{x}_{pre} be a finite walk and \mathbf{x}_{suf} be a finite cycle. A trajectory \mathbf{x} is in *prefix-suffix* form if it is of the form $\mathbf{x} = \mathbf{x}_{\text{pre}}(\mathbf{x}_{\text{suf}})^\omega$, where ω denotes infinite repetition.

We will require that the (time-varying) labeling function L_t is eventually periodic.

Assumption 1 There exists a finite $t' \in \mathcal{T}^\infty$ and a $\Omega \in \mathbb{N}$ such that $L_t = L_{t+\Omega}$ for all $t \geq t' \in \mathcal{T}^\infty$. We further assume that Ω is minimal among all possible values.

In the sequel, we will only consider trajectories $\mathbf{x} = \mathbf{x}_{\text{pre}}(\mathbf{x}_{\text{suf}})^\omega$ in prefix-suffix form. While both \mathbf{x}_{pre} and \mathbf{x}_{suf} are finite, the constraint that \mathbf{x}_{suf} is a cycle allows us to repeat that sequence of states forever. Repeating the same sequence of states is a sufficient condition that the word $L(\mathbf{x}_{\text{suf}})$ (i.e., the sequence of atomic propositions) is also repeated (using Assumption 1). However, only the word matters for the feasibility of an LTL formula, not the exact sequence of states. In fact, there may exist other trajectories that produce the same word $L(x)$, but are not eventually periodic. Our approach cannot find such trajectories, although we have not noticed this limitation in our experiments. This differs from the case of finite discrete systems, where a prefix-suffix form is sufficient to find a feasible solution if one exists [2].

In the next section, we will encode the temporal operators as mixed-integer constraints on \mathbf{x}_{pre} and \mathbf{x}_{suf} . Let $\mathbf{x}_{\text{cat}} := \mathbf{x}_{\text{pre}}\mathbf{x}_{\text{suf}}$ denote the concatenation of \mathbf{x}_{pre} and \mathbf{x}_{suf} , and assign time indices to \mathbf{x}_{cat} as $\mathcal{T}_{\text{cat}} := \{0, 1, \dots, T_s, \dots, T\}$. Let $\mathcal{T}_{\text{pre}} := \{0, 1, \dots, T_s - 1\}$ and $\mathcal{T}_{\text{suf}} := \{T_s, \dots, T\}$, where T_s is the first time instance on the suffix. The infinite repetition of \mathbf{x}_{suf} is enforced by the constraint $\mathbf{x}_{\text{cat}}(T_s) = f(\mathbf{x}_{\text{cat}}(T), u)$ for some $u \in \mathcal{U}$. By Assumption 1, it is sufficient that T_s is greater

than t' and that the length of \mathcal{T}_{suf} is an integer multiple of Ω . We often identify $\mathbf{x}_{\text{pre}}(0) \cdots \mathbf{x}_{\text{pre}}(T_{\text{pre}})$ with $\mathbf{x}_{\text{cat}}(0) \cdots \mathbf{x}_{\text{cat}}(T_s - 1)$ and $\mathbf{x}_{\text{suf}}(0) \cdots \mathbf{x}_{\text{suf}}(T_{\text{suf}})$ with $\mathbf{x}_{\text{cat}}(T_s) \cdots \mathbf{x}_{\text{cat}}(T)$ in the obvious manner.

5 A Mixed-Integer Linear Formulation of LTL Constraints

In this section, we develop a mixed-integer programming formulation for a given prefix-suffix trajectory parameterization, $\mathbf{x}_{\text{cat}} = \mathbf{x}_{\text{pre}}\mathbf{x}_{\text{suf}}$. The corresponding system trajectory is $\mathbf{x} = \mathbf{x}_{\text{pre}}(\mathbf{x}_{\text{suf}})^\omega$. Since the system is deterministic, this defines a corresponding control input sequence. The split between \mathbf{x}_{pre} and \mathbf{x}_{suf} can either be specified *a priori* or left as a variable (see [26] for details). We mix notation in the following and refer to x and \mathcal{T} instead of \mathbf{x}_{cat} and \mathcal{T}_{cat} when clear from context.

5.1 Relating the Dynamics and Propositions

We now relate the state of a system to the set of atomic propositions that are *True* at each time instance. We assume that each propositional formula ψ is described at time t by the union of a finite number of polytopes, indexed by the finite index set I_t^ψ . Let $\llbracket \psi \rrbracket(t) := \{x \in \mathcal{X} \mid H_t^{\psi_i} x \leq K_t^{\psi_i} \text{ for some } i \in I_t^\psi\}$ represent the set of states that satisfy propositional formula ψ at time t . We assume that these have been constructed as necessary from the system's original atomic propositions. We note that a proposition preserving partition [1] is not necessary or even desired.

For each propositional formula ψ , introduce binary variables $z_t^{\psi_i} \in \{0, 1\}$ for all $i \in I_t^\psi$ and for all $t \in \mathcal{T}$. Let x_t be the state of the system at time t and M be a vector of sufficiently large constants. The *big-M* formulation

$$\begin{aligned} H_t^{\psi_i} x_t &\leq K_t^{\psi_i} + M(1 - z_t^{\psi_i}), \quad \forall i \in I_t^\psi \\ \sum_{i \in I_t^\psi} z_t^{\psi_i} &= 1 \end{aligned} \tag{4}$$

enforces the constraint that $x_t \in \llbracket \psi \rrbracket(t)$ at time t . Define $P_t^\psi := \sum_{i \in I_t^\psi} z_t^{\psi_i}$. If $P_t^\psi = 1$, then $x_t \in \llbracket \psi \rrbracket(t)$. If $P_t^\psi = 0$, then nothing can be inferred.

The big-M formulation may give poor continuous relaxations of the binary variables, i.e., $z_t^{\psi_i} \in [0, 1]$, which may lead to poor performance during optimization [23]. Such relaxations are frequently used during the solution of mixed-integer linear programs [23]. Thus, we introduce an alternate representation whose continuous relaxation is the convex hull of the original set $\llbracket \psi \rrbracket(t)$. This formulation is well-known in the optimization community [12], but does not appear in the trajectory

generation literature ([8, 20, 24] and references therein). As such, this formulation may be of independent interest for trajectory planning with obstacles.

The *convex hull* formulation

$$\begin{aligned}
 H_t^{\psi_i} x_t^i &\leq K_t^{\psi_i} z_t^{\psi_i}, \quad \forall i \in I_t^{\psi} \\
 \sum_{i \in I_t^{\psi}} z_t^{\psi_i} &= 1, \\
 \sum_{i \in I_t^{\psi}} x_t^i &= x_t
 \end{aligned} \tag{5}$$

represents the same set as the big-M formulation (4). While the convex hull formulation introduces more continuous variables, it gives the tightest linear relaxation of the disjunction of the polytopes and reduces the need to select the M parameters [12]. Note that we will only use the convex hull formulation (5) for safety and persistence formulas (i.e., φ_{safe} and φ_{per}) in Sect. 5.2, as $P_t^{\psi} = 0$ enforces $x = 0$.

Regardless if one uses the big-M or convex hull formulation, only one binary variable is needed for each polyhedron (i.e., finite conjunction of halfspaces). This compares favorably with the approach in [15], where a binary variable is introduced for each halfspace. Additionally, the auxiliary continuous variables and mixed-integer constraints previously used are not needed because we use implication. For simple tasks such as $\varphi = \Diamond\psi$, our method can use significantly fewer binary variables than previously needed, depending on the number of halfspaces and polytopes needed to describe $\llbracket \psi \rrbracket$.

For every temporal operator described in the following section, the constraints in (4) or (5) should be understood to be implicitly applied to the corresponding propositional formulas so that $P_t^{\psi} = 1$ implies that the system satisfies ψ at time t . Also, note that we use different binary variables for each formula—even when representing the same set.

5.2 The Mixed-Integer Linear Constraints

In this section, the trajectory parameterization \mathbf{x} has been *a priori* split into a prefix \mathbf{x}_{pre} and a suffix \mathbf{x}_{suf} . This assumption can be relaxed, so that the size of \mathbf{x}_{pre} and \mathbf{x}_{suf} are optimization variables (see [26] for details). We further assume that \mathbf{x}_{pre} and \mathbf{x}_{suf} satisfy Assumption 1.

In the following, the correctness of the constraints applied to \mathbf{x}_{pre} and \mathbf{x}_{suf} comes directly from the temporal logic semantics given in Sect. 2.2 and the form of the trajectory $\mathbf{x} = \mathbf{x}_{\text{pre}}(\mathbf{x}_{\text{suf}})^{\omega}$. The most important factors are whether a property can be verified over finite- or infinite-horizons. All infinite-horizon (liveness) properties must be satisfied on the suffix \mathbf{x}_{suf} .

We begin with the fundamental temporal operators Φ_{core} . Safety and persistence require a mixed-integer linear constraint for each time step, while guarantee and liveness only require a single mixed-integer linear constraint.

Safety, $\varphi_{\text{safe}} = \Box\psi$, is satisfied by the constraints

$$\begin{aligned} P_t^\psi &= 1, \quad \forall t \in \mathcal{T}_{\text{pre}}, \\ P_t^\psi &= 1, \quad \forall t \in \mathcal{T}_{\text{suf}}, \end{aligned}$$

which ensure that the system is always in a $\llbracket\psi\rrbracket$ region. Similarly, persistence, $\varphi_{\text{per}} = \Diamond\Box\psi$, is enforced by

$$P_t^\psi = 1, \quad \forall t \in \mathcal{T}_{\text{suf}},$$

which ensures the system eventually remains in a $\llbracket\psi\rrbracket$ region.

Guarantee, $\varphi_{\text{goal}} = \Diamond\psi$, is satisfied by the constraints

$$\sum_{t \in \mathcal{T}_{\text{pre}}} P_t^\psi + \sum_{t \in \mathcal{T}_{\text{suf}}} P_t^\psi = 1,$$

which ensures the system eventually visits a $\llbracket\psi\rrbracket$ region. Similarly, liveness $\varphi_{\text{live}} = \Box\Diamond\psi$ is enforced by

$$\sum_{t \in \mathcal{T}_{\text{suf}}} P_t^\psi = 1,$$

which ensures the system repeatedly visits a $\llbracket\psi\rrbracket$ region.

Until, $\varphi_{\text{until}} = \psi \mathcal{U} \phi$, is enforced by

$$\begin{aligned} P_0^\phi &= s_0, \\ P_t^\phi &= s_t - s_{t-1}, \quad t = 1, \dots, T \\ P_t^\psi &= 1 - s_t, \quad \forall t \in \mathcal{T}, \end{aligned}$$

where we use auxiliary binary variables $s_t \in \{0, 1\}$ for all $t \in \mathcal{T}$ such that $s_t \leq s_{t+1}$ for $t = 0, \dots, T-1$ and $s_T = 1$.

Now consider the response temporal operators Φ_{resp} . For these formulas, the definition of implication is used to convert each inner formula into a disjunction between a property that holds at a state and a property that holds at some point in the future. The response formulas require a mixed-integer linear constraint for each time step.

For next-step response, $\varphi_{\text{resp}}^1 = \Box(\psi \implies \bigcirc\phi) = \Box(\neg\psi \vee \bigcirc\phi)$, the additional constraints are

$$\begin{aligned} P_t^{\neg\psi} + P_{t+1}^{\phi} &= 1, \quad t = 0, \dots, T_s, \dots, T-1, \\ P_T^{\neg\psi} + P_{T_s}^{\phi} &= 1, \end{aligned}$$

Similarly, steady-state next-step response, $\varphi_{\text{resp}}^3 = \Diamond\Box(\psi \implies \bigcirc\phi) = \Diamond\Box(\neg\psi \vee \bigcirc\phi)$, is satisfied by

$$\begin{aligned} P_t^{\neg\psi} + P_{t+1}^{\phi} &= 1, \quad t = T_s, \dots, T-1, \\ P_T^{\neg\psi} + P_{T_s}^{\phi} &= 1, \end{aligned}$$

Eventual response, $\varphi_{\text{resp}}^2 = \Box(\psi \implies \Diamond\phi) = \Box(\neg\psi \vee \Diamond\phi)$, requires the following constraints

$$\begin{aligned} P_t^{\neg\psi} + \sum_{\tau=t}^T P_{\tau}^{\phi} &= 1, \quad \forall t \in \mathcal{T}_{\text{pre}}, \\ P_t^{\neg\psi} + \sum_{t \in \mathcal{T}_{\text{suf}}} P_t^{\phi} &= 1, \quad \forall t \in \mathcal{T}_{\text{suf}}. \end{aligned}$$

Similarly, for steady-state eventual response, $\varphi_{\text{resp}}^4 = \Diamond\Box(\psi \implies \Diamond\phi) = \Diamond\Box(\neg\psi \vee \Diamond\phi)$, the additional constraints are

$$P_t^{\neg\psi} + \sum_{t \in \mathcal{T}_{\text{suf}}} P_t^{\phi} = 1, \quad \forall t \in \mathcal{T}_{\text{suf}}.$$

Now consider the fairness temporal operators Φ_{fair} . In the following, the definition of implication is used to rewrite the inner formula as disjunction between a single safety (persistence) property and a conjunction of guarantee (liveness) properties. These formulas require a mixed-integer linear constraint for each conjunction in the response and each time step.

Conditional goal visitation, $\varphi_{\text{fair}}^1 = \Diamond\psi \implies \bigwedge_{j=1}^m \Diamond\phi_j = \Box\neg\psi \vee \bigwedge_{j=1}^m \Diamond\phi_j$, is specified by

$$P_t^{\neg\psi} + \sum_{t \in \mathcal{T}} P_t^{\phi_j} = 1, \quad \forall j = 1, \dots, m, \forall t \in \mathcal{T}.$$

Conditional repeated goal visitation, $\varphi_{\text{fair}}^2 = \Diamond\psi \implies \bigwedge_{j=1}^m \Box\Diamond\phi_j = \Box\neg\psi \vee \bigwedge_{j=1}^m \Box\Diamond\phi_j$, is enforced as

$$P_t^{\neg\psi} + \sum_{t \in \mathcal{T}_{\text{suf}}} P_t^{\phi_j} = 1, \quad \forall j = 1, \dots, m, \forall t \in \mathcal{T}.$$

Similarly, $\varphi_{\text{fair}}^3 = \Box\Diamond\psi \implies \bigwedge_{j=1}^m \Box\Diamond\phi_j = \Diamond\Box\neg\psi \vee \bigwedge_{j=1}^m \Box\Diamond\phi_j$, is represented by

$$P_t^{\neg\psi} + \sum_{t \in \mathcal{T}_{\text{suf}}} P_t^{\phi_j} = 1, \quad \forall j = 1, \dots, m, \quad \forall t \in \mathcal{T}_{\text{suf}}.$$

We have encoded the temporal logic specifications on the system variables using mixed-integer linear constraints. Note that the equality constraints on the binary variables dramatically reduce search space. In Sect. 6 we discuss adding dynamics to further constrain the possible behaviors of the system.

6 System Dynamics

The mixed-integer constraints in Sect. 5 are over a sequence of states, and thus are independent of the specific system dynamics. Dynamic constraints on the sequence of states can also be enforced by standard transcription methods [5]. However, the resulting optimization problem may then be a mixed-integer *nonlinear* program due to the dynamics. We highlight two useful classes of nonlinear systems where the dynamics can be encoded using mixed-integer *linear* constraints.

6.1 Mixed Logical Dynamical Systems

Mixed Logical Dynamical (MLD) systems have both continuous and discrete-valued states and allow one to model nonlinearities, logic, and constraints [4]. These systems include constrained linear systems, linear hybrid automata, and piecewise affine systems. An MLD system is of the form

$$\begin{aligned} x(t+1) &= Ax(t) + B_1u(t) + B_2\delta(t) + B_3z(t) \\ \text{subject to } E_2\delta(t) + E_3z(t) &\leq E_1u(t) + E_4x(t) + E_5, \end{aligned}$$

where $t \in \mathcal{T}^\infty$, $x \in \mathcal{X} \subseteq \mathbb{R}^{n_c} \times \{0, 1\}^{n_l}$ are the continuous and binary states, $u \in \mathcal{U} \subseteq \mathbb{R}^{m_c} \times \{0, 1\}^{m_l}$ are the inputs, and $\delta \in \{0, 1\}^{r_l}$ and $z \in \mathbb{R}^{r_l}$ are auxiliary binary and continuous variables, respectively. The terms A , B_1 , B_2 , B_3 , E_1 , E_2 , E_3 , E_4 , and E_5 are system matrices of appropriate dimension. We assume that the system is deterministic and well-posed (see Definition 1 in [4]).

6.2 Differentially Flat Systems

A system is *differentially flat* if there exists a set of outputs such that all states and control inputs can be determined from these outputs without integration. If a system has states $x \in \mathbb{R}^n$ and control inputs $u \in \mathbb{R}^m$, then it is flat if we can find

outputs $y \in \mathbb{R}^m$ of the form $y = y(x, u, \dot{u}, \dots, u^{(p)})$ such that $x = x(y, \dot{y}, \dots, y^{(q)})$ and $u = u(y, \dot{y}, \dots, y^{(q)})$. Thus, we can plan trajectories in output space and then map these to control inputs [17].

Differentially flat systems may be encoded using mixed integer linear constraints in certain cases, e.g., the flat output is constrained by mixed integer linear constraints. This holds for relevant classes of robotic systems, including quadrotors and car-like robots. However, control input constraints are typically non-convex in the flat output. Common approaches to satisfy control constraints are to plan a sufficiently smooth trajectory or slow down along a trajectory [19].

7 Examples

We demonstrate our techniques on a variety of motion planning problems. The first example is a chain of integrators parameterized by dimension. Our second example is a quadrotor model from [25]. Our final example is a nonlinear car-like vehicle with drift. All computations were done on a laptop with a 2.4 GHz dual-core processor and 4 GB of memory using CPLEX [23] through Yalmip [18].

The environment and task is motivated by a pickup and delivery scenario. All properties should be understood to be with respect to regions in the plane (see Fig. 1). Let P be a region where supplies can be picked up and $D1$ and $D2$ be regions where supplies must be delivered. The robot must remain in the safe region S (in white). Formally, the task specification is $\varphi = \Box S \wedge \Box \Diamond P \wedge \Box \Diamond D1 \wedge \Box \Diamond D2$. Additionally, we minimize the maximum cost function (3) where $c(x_t, u_t) = |u_t|$ penalizes the control input.

In the remainder of this section, we consider this temporal logic motion planning problem for different system models. We use the simultaneous (sim.) approach described in Sect. 5.2, and also a sequential (seq.) approach from [26] that first computes the suffix and then the prefix. A trajectory of length 60 (split evenly between the prefix and suffix) is used in all cases, and all results are averaged over 20 randomly generated environments. The simultaneous approach uses between 300 and 469 binary variables with a mean of 394. Finally, all continuous-time models are discretized using a first-order hold and time-step of 0.5 s.

7.1 Chain of Integrators

The first system is a chain of orthogonal integrators in the x and y directions. The k th derivative of the x and y positions are controlled, i.e., $x^{(k)} = u_x$ and $y^{(k)} = u_y$, subject to the constraints $|u_x| \leq 0.5$ and $|u_y| \leq 0.5$. The general state constraints are $|x^{(i)}| \leq 1$ and $|y^{(i)}| \leq 1$ for $i = 1, \dots, k-1$. Results are given in Tables 1 and 2 under “chain-2”, “chain-6”, and “chain-10”, where “chain- k ” indicates that the k th derivative in both the x and y positions is controlled.

Table 1 Time until a feasible solution was found (mean \pm standard error) and number of problems (out of 20) solved in 45 s using the big-M formulation (4) with $M = 10$

Model	Dimensional	Feasible solution (s)		Number of problems solved	
		Simultaneous	Sequential	Simultaneous	Sequential
Chain-2	4	1.10 ± 0.09	0.64 ± 0.06	20	20
Chain-6	12	$4.70 \pm .48$	2.23 ± 0.15	20	20
Chain-10	20	9.38 ± 1.6	3.74 ± 0.29	20	19
Quadrotor	10	4.20 ± 0.66	1.80 ± 0.15	20	20
Quadrotor-flat	10	2.26 ± 0.36	1.99 ± 1.0	20	20
Car-3	3	43.9 ± 0.77	10.7 ± 2.0	4	20
Car-4	3	42.4 ± 1.7	18.7 ± 3.1	2	18
Car-flat	3	15.8 ± 3.8	14.0 ± 4.4	12	14

Table 2 Time until a feasible solution was found (mean \pm standard error) and number of problems (out of 20) solved in 45 s using the convex hull formulation (5)

Model	Dimensional	Feasible solution (s)		Number of problems solved	
		Simultaneous	Sequential	Simultaneous	Sequential
Chain-2	4	1.94 ± 0.23	0.94 ± 0.11	20	20
Chain-6	12	12.4 ± 2.7	2.89 ± 0.32	20	20
Chain-10	20	16.9 ± 3.0	7.28 ± 1.2	17	15
Quadrotor	10	18.9 ± 3.8	2.80 ± 0.35	16	20
Car-3	3	37.3 ± 3.1	13.3 ± 1.6	8	20

7.2 Quadrotor

We now consider the quadrotor model used in [25] for point-to-point motion planning, to which we refer the reader for a complete description of the model. The state $x = (p, v, r, w)$ is 10-dimensional, consisting of position $p \in \mathbb{R}^3$, velocity $v \in \mathbb{R}^3$, orientation $r \in \mathbb{R}^2$, and angular velocity $w \in \mathbb{R}^2$. This model is the linearization of a nonlinear model about hover with the yaw constrained to be zero. The control input $u \in \mathbb{R}^3$ is the total, roll, and pitch thrust. Results are given in Tables 1 and 2 under “quadrotor”, and a sample trajectory is shown in Fig. 1.

Also, we use the fact that the quadrotor is differentially flat [19] to generate trajectories for the nonlinear model (with fixed yaw). We parameterize the flat output $p \in \mathbb{R}^3$ with eight piecewise polynomials of degree three, and then optimize over their coefficients to compute a smooth trajectory. Afterwards, we check that the trajectory does not violate the control input constraints. Results are given in Table 1 under “quadrotor-flat”.

7.3 Nonlinear Car

Consider a nonlinear car-like vehicle with state $x = (p_x, p_y, \theta)$ and dynamics $\dot{x} = (v \cos(\theta), v \sin(\theta), u)$. The variables p_x, p_y are position (m) and θ is orientation (rad). The vehicle's speed v is fixed at 0.8(m/s) and its control input is constrained as $|u| \leq 2.5$. We form a hybrid MLD model by linearizing the system about different orientations $\hat{\theta}_i$ for $i = 1, \dots, k$. The dynamics are governed by the closest linearization to the current θ . Results with $k = 3$ and $k = 4$ are given in Table 1 under “car-3” and “car-4”, respectively. A sample trajectory of “car-4” is show in Fig. 1.

Additionally, we use the flat output $(x, y) \in \mathbb{R}^2$ to generate trajectories for the nonlinear car-like model in a similar manner as for the quadrotor model. Results are given in Table 1 under “car-flat”.

7.4 Discussion and Comparison

We first compare our approach to reachability-based algorithms that compute a finite abstraction [16, 28]. We used the method in [28] to compute a discrete abstraction for a two dimensional system in 22 s, and [16] reports abstracting a four dimensional system in just over a minute. This contrasts with our mixed-integer approach that can routinely find solutions to such problems in seconds, although we do not compute a feedback controller. Our results appear particularly promising for situations where the environment is dynamically changing and a finite abstraction must be repeatedly computed.

We also compare to the finite-horizon mixed-integer formulation given in [15]. Consider the task $\varphi = \Diamond \psi$, where $\llbracket \psi \rrbracket$ is a convex polytope defined by m halfspaces. Our method uses one binary variable at each time step, while their approach uses m . Additionally, while we encode eventually (\Diamond) using a single constraint, their approach uses a number of constraints quadratic in the trajectory length.

In most of our examples, we are able to quickly compute a feasible trajectory that satisfies a temporal logic formula by solving a mixed-integer linear program. This is aided by the sequential approach, which separates the problem into computing a suffix and then a prefix [26]. It typically takes a long time to compute a trajectory that is provably globally optimal, although this does happen in finite time.

Finally, the convex hull formulation performed poorly in our examples. There is an empirical tradeoff between having tighter continuous relaxations and the number of continuous variables in the formulation. We hypothesize that the convex hull formulation will be most useful in cases when (1) the number of binary variables is large, or (2) the cost function is minimized near the boundary of the region.

8 Conclusion

We presented a novel mixed-integer programming-based method for control of nonlinear systems with a useful fragment of LTL that allows both finite- and infinite-horizon properties to be specified. Our method is efficient in the number of binary variables used to model the an LTL formula. Additionally, we showed the computational effectiveness of our approach on temporal logic motion planning examples.

Future work will consider reactive environments by including both continuous and discrete disturbances using a receding horizon control approach. Additionally, we will expand the space of tasks that can be specified by including additional temporal operators and timing constraints.

Acknowledgments The authors thank Matanya Horowitz, Scott Livingston, and Ufuk Topcu for helpful feedback. This work was supported by a NDSEG Fellowship and the Boeing Corporation.

References

1. Alur, R., Henzinger, T.A., Lafferriere, G., Pappas, G.J.: Discrete abstractions of hybrid systems. *Proc. IEEE* **88**(7), 971–984 (2000)
2. Baier, C., Katoen, J.P.: *Principles of Model Checking*. MIT Press, Cambridge (2008)
3. Belta, C., Habets, L.C.G.J.M.: Controlling of a class of nonlinear systems on rectangles. *IEEE Trans. Autom. Control* **51**, 1749–1759 (2006)
4. Bemporad, A., Morari, M.: Control of systems integrating logic, dynamics, and constraints. *Automatica* **35**, 407–427 (1999)
5. Betts, J.T.: *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd edition. SIAM (2000)
6. Bhatia, A., Maly, M.R., Kavraki, L.E., Vardi, M.Y.: Motion planning with complex goals. *IEEE Robot. Autom. Mag.* **18**, 55–64 (2011)
7. Blair, C.E., Jeroslow, R.G., Lowe, J.K.: Some results and experiments in programming techniques for propositional logic. *Comput. Oper. Res.* **13**, 633–645 (1986)
8. Earl, M.G., D’Andrea, R.: Iterative MILP methods for vehicle-control problems. *IEEE Trans. Robot.* **21**, 1158–1167 (2005)
9. Fainekos, G.E., Girard, A., Kress-Gazit, H., Pappas, G.J.: Temporal logic motion planning for dynamic robots. *Automatica* **45**, 343–352 (2009)
10. Habets, L., Collins, P.J., van Schuppen, J.H.: Reachability and control synthesis for piecewise-affine hybrid systems on simplices. *IEEE Trans. Autom. Control* **51**, 938–948 (2006)
11. Hooker, J.N., Fedjki, C.: Branch-and-cut solution of inference problems in propositional logic. *Ann. Math. Artif. Intell.* **1**, 123–139 (1990)
12. Jeroslow, R.G.: Representability in mixed integer programming, I: Characterization results. *Disc. Appl. Math.* **17**, 223–243 (1987)
13. Karaman, S., Frazzoli, E.: Linear temporal logic vehicle routing with applications to multi-UAV mission planning. *Int. J. Robust Nonlinear Control* **21**, 1372–1395 (2011)
14. Karaman, S., Frazzoli, E.: Sampling-based algorithms for optimal motion planning with deterministic μ -calculus specifications. In: *Proceedings of the American Control Conference* (2012)
15. Karaman, S., Sanfelice, R.G., Frazzoli, E.: Optimal control of mixed logical dynamical systems with linear temporal logic specifications. In: *Proceedings of the IEEE Conference on Decision and Control*, pp. 2117–2122 (2008)
16. Kloetzer, M., Belta, C.: A fully automated framework for control of linear systems from temporal logic specifications. *IEEE Trans. Autom. Control* **53**(1), 287–297 (2008)

17. LaValle, S.M.: Planning Algorithms. Cambridge University Press, Cambridge (2006)
18. Löfberg, J.: YALMIP: A toolbox for modeling and optimization in MATLAB. In: Proceedings of the CACSD Conference. Taipei, Taiwan (2004). Software available at <http://control.ee.ethz.ch/~joloef/yalmip.php>
19. Mellinger, D., Kushleyev, A., Kumar, V.: Mixed-integer quadratic program trajectory generation for heterogeneous quadrotor teams. In: Proceedings of the International Conference on Robotics and Automation (2012)
20. Richards, A., How, J.P.: Aircraft trajectory planning with collision avoidance using mixed integer linear programming. In: American Control Conference (2002)
21. Smith, S.L., Tumova, J., Belta, C., Rus, D.: Optimal path planning for surveillance with temporal-logic constraints. *Int. J. Robot. Res.* **30**, 1695–1708 (2011)
22. Toth, P., Vigo, D. (eds.): The Vehicle Routing Problem. Philadelphia, PA: SIAM (2001)
23. User's Manual for CPLEX V12.2. IBM (2010)
24. Vitus, M.P., Pradeep, V., Hoffmann, J., Waslander, S.L., Tomlin, C.J.: Tunnel-MILP: path planning with sequential convex polytopes. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference (2008)
25. Webb, D.J., van den Berg, J.: Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In: Proceedings of the IEEE International Conference on Robotics and Automation (2013)
26. Wolff, E.M., Murray, R.M.: Optimal control of mixed logical dynamical systems with long-term temporal logic specifications. Technical report, California Institute of Technology (2013). <http://resolver.caltech.edu/CaltechCDSTR:2013.001>
27. Wolff, E.M., Topcu, U., Murray, R.M.: Optimal control with weighted average costs and temporal logic specifications. In: Proceedings of the Robotics: Science and Systems (2012)
28. Wongpiromsarn, T., Topcu, U., Ozay, N., Xu, H., Murray, R.M.: TuLiP: A software toolbox for receding horizon temporal logic planning. In: Proceedings of the International Conference on Hybrid Systems: Computation and Control (2011). <http://tulip-control.sf.net>
29. Wongpiromsarn, T., Topcu, U., Murray, R.M.: Receding horizon temporal logic planning. *IEEE Trans. Autom. Control* (2012)

Robotics Research

The 16th International Symposium ISRR

Inaba, M.; Corke, P. (Eds.)

2016, XIV, 701 p. 298 illus., 33 illus. in color., Hardcover

ISBN: 978-3-319-28870-3