UC Berkeley
UC Berkeley Previously Published Works

Title
Reactive synthesis from signal temporal logic specifications

Permalink
https://escholarship.org/uc/item/6b86d7vh

ISBN
9781450334334

Authors
Raman, Vasumathi
Donzé, Alexandre
Sadigh, Dorsa
et al.

Publication Date
2015-04-14

DOI
10.1145/2728606.2728628

Peer reviewed

Reactive Synthesis from Signal Temporal Logic Specifications

Vasumathi Raman
California Institute of
Technology
Pasadena, CA, USA
vasu@caltech.edu

Alexandre Donzé
UC Berkeley
Berkeley, CA, USA
donze@berkeley.edu

Dorsa Sadigh
UC Berkeley
Berkeley, CA, USA
dsadigh@berkeley.edu

Richard M. Murray
California Institute of
Technology
Pasadena, CA, USA
murray@cds.caltech.edu

Sanjit A. Seshia
UC Berkeley
Berkeley, CA, USA
sseshia@eecs.berkeley.edu

ABSTRACT

We present a counterexample-guided inductive synthesis approach to controller synthesis for cyber-physical systems subject to signal temporal logic (STL) specifications, operating in potentially adversarial nondeterministic environments. We encode STL specifications as mixed integer-linear constraints on the variables of a discrete-time model of the system and environment dynamics, and solve a series of optimization problems to yield a satisfying control sequence. We demonstrate how the scheme can be used in a receding horizon fashion to fulfill properties over unbounded horizons, and present experimental results for reactive controller synthesis for case studies in building climate control and autonomous driving.

1. INTRODUCTION

We are concerned with controlling hybrid systems to satisfy desired properties despite a potentially adversarial environment; the provided solution must be robust to environment actions with regards to which we are uncertain. Recently, temporal logics have proven a valuable tool for controller synthesis, because they provide a compact mathematical formalism for specifying desired behaviors of a system. There is a rich body of literature containing algorithms for verification and synthesis of systems obeying temporal logic specifications. Approaches can be broadly categorized based on whether they utilize a discrete abstraction of the system, and whether the environment is assumed to be deterministic.

Approaches that utilize a discrete abstraction enable construction of discrete supervisory controllers, which have successfully been used to construct hybrid controllers for domains including robotics and aircraft power system design; these include approaches that deal with deterministic [17, 23] as well as adversarial environments [10, 27]. In contrast, approaches that eschew discrete abstractions include those based on sampling-based methods [15], and mixed-integer linear programming encodings of temporal logic specifications[16, 14, 18, 26, 24]. The latter have thus far been confined to the realm of deterministic operating environments, and it is this gap that we close with the current work.

We adopt a *counterexample-guided inductive synthesis* [25] approach to synthesize a controller satisfying reactive specifications. Inductive synthesis refers to the automated generation of a system from input-output examples, using each new example to iteratively refine the hypothesis about the system until convergence. In Counterexample-Guided Inductive Synthesis (CEGIS), the examples are mostly counterexamples discovered while trying to verify correctness of the current guess. CEGIS thus relies primarily on a validation engine to validate candidates produced at intermediate iterations, which can produce counterexamples for use in the next iteration. Automated synthesis of systems using CEGIS and the closely related Counterexample-Guided Abstraction Refinement (CEGAR) paradigm has been widely studied in various contexts [6, 2, 13].

The specification language adopted here is Signal Temporal Logic (STL) [21], which allows the specification of temporal properties of real-valued signals, and has been applied to the analysis of hybrid dynamical systems from various application domains such as analog and mixed signal circuits, systems biology or Cyber-Physical Systems (CPS). STL has the advantage of naturally admitting a *quantitative* semantics which, in addition to the binary answer to the question of satisfaction, provides a real number indicating the quality of the satisfaction or violation. Such quantitative semantics have been defined for timed logics e.g. Metric Temporal Logic (MTL) [11] and STL [8] to assess the *robustness* of the systems to parameter or timing variations. We exploit this ability to compute the robustness of satisfaction in the validation engine for our CEGIS approach to reactive synthesis.

A key advantage of temporal logic over, e.g., domain-specific

languages based on propositional logic, is that it allows the expression of properties of infinite traces. We would therefore like to synthesize controllers that can run indefinitely, and satisfy infinite-horizon properties. Receding Horizon Control (RHC) [22] is based on iterative, finite horizon, discrete time optimization of a model of the plant: at time t, the current plant state is observed, and an optimal control strategy is computed for a finite time horizon in the future, $[t, t + H]$. The first step of the computed strategy is implemented, the plant state is then sampled again, and new calculations performed on a horizon of size H starting from the new current state. This not only reduces computational complexity, but improves robustness with respect to exogenous disturbances and modeling uncertainties by allowing new information to be incorporated as it becomes available [22].

We have already made the connection between Receding Horizon Control (RHC) and control synthesis from STL specifications in previous work[24], where we specify desired properties of the system using a STL formula, and synthesize control such that the system satisfies that specification, while using a receding horizon approach. We presented automatically-generated Mixed Integer Linear Program (MILP) encodings for STL specifications, extending the Bounded Model Checking (BMC) paradigm for finite discrete systems [5] to STL. These encodings can be used not only to generate open-loop control signals that satisfy finite and infinite horizon STL properties, but also to generate signals that maximize quantitative (robust) satisfaction. In this paper, we show how the robustness-based encoding can be used to produce a validation engine that synthesizes counterexamples to guide a CEGIS approach to reactive synthesis.

Abbas et al. [1] exploit the quantitative semantics of Metric Temporal Logic (MTL) to design a framework for specification-guided testing of stochastic cyber-physical systems. Leveraging results from stochastic optimization, they frame the verification of properties on such systems as a global optimization problem of minimizing the expected robustness of satisfaction. While we deal with nondeterministic systems rather than stochastic systems, our CEGIS scheme uses a similar idea when finding an adversarial environment input that minimizes the robustness of satisfaction. Bartocci et al. [3] also apply the definition of robustness to a stochastic model and formulate an optimization problem over system parameters in order to maximize the average robustness of satisfying a temporal logic formula. Their work can be viewed as synthesis based on robustness for system parameters, but involves approximating the distribution of the robustness score for the stochastic system being considered.

Receding horizon control that satisfies temporal logic specifications in adversarial settings has been considered before in the context of Linear Temporal Logic (LTL) [27], where the authors propose a scheme that makes use of discrete abstractions to synthesize supervisory controllers for specifications in the GR(1) subset of LTL. In that work, feasibility of the global specification is determined via symbolic checks on a series of pre-defined smaller problems, and strategies extracted as needed. In contrast, we do not require an *a priori* defined finite set of sub-problems. Our approach also ex-

tends synthesis capabilities to a wider class of temporal logic specifications and environments than [12, 4], and avoids potentially expensive computations of a finite state abstraction of the system as in [7] and [27].

Contributions: The key novel contribution of this paper is a CEGIS approach to controller synthesis for cyber-physical systems subject to signal temporal logic (STL) specifications, operating in potentially adversarial nondeterministic environments. Specific features of our approach include:

- We leverage our previously-proposed encoding of STL specifications as mixed integer-linear constraints on the variables of a discrete-time model of the system and environment dynamics [24], and solve a counterexample-guided series of optimization problems to yield a satisfying control sequence.

- Our scheme can be used in a receding horizon fashion to fulfill properties over unbounded horizons.

- We present experimental results using a case study of controller synthesis on a model of a Heating Ventilation and Air Conditioning (HVAC) system with nondeterministic elements in the environment, and an autonomous driving scenario in the presence of adversarial agents; simulation results in these two domains illustrate the effectiveness of our methodology.

Our method is a fundamentally novel approach to *reactive* synthesis for hybrid systems, different from most current methods, which often rely on model transformations (e.g., abstraction and discretization).

2. PRELIMINARIES

We consider a continuous-time system Σ of the form

$$\dot{x} = f(x, u, w)$$

where $x \in \mathcal{X} \subseteq (\mathbb{R}^{n_c} \times \{0, 1\}^{n_l})$ are the continuous and binary/logical states, $u \in U \subseteq (\mathbb{R}^{m_c} \times \{0, 1\}^{m_l})$ are the (continuous and logical) control inputs, $w \in W \subseteq (\mathbb{R}^{e_c} \times \{0, 1\}^{e_l})$ are the (possibly adversarial) external inputs or disturbances, and $x_0 \in \mathcal{X}$ is the initial state. We will refer to w as the *environment* input.

Given a sampling time $\Delta t > 0$, we assume that Σ admits a discrete-time approximation Σ_d of the form

$$x(t_{k+1}) = f_d(x(t_k), u(t_k), w(t_k)) \qquad (1)$$

where for all $k > 0$, $t_{k+1} - t_k = \Delta t$. A *run*

$$\xi = (x_0 u_0 w_0)(x_1 u_1 w_1)(x_2 u_2 w_2)...$$

of Σ_d is a sequence where $x_k = x(t_k) \in \mathcal{X}$ is the state of the system at index k, and for each $k \in \mathbb{N}$, $u_k = u(t_k) \in U$, $w_k = w(t_k) \in W$ and $x_{k+1} = f_d(x_k, u_k, w_k)$. Given $x_0 \in X$, $\mathbf{u} \in U^\omega$ and $\mathbf{w} \in W^\omega$, denote by $\xi(x_0, \mathbf{u}, \mathbf{w})$ the run generated following equation (1). The corresponding sequence of states, which we also call the *discrete-time signal*, or simply *signal*, is denoted by $\mathbf{x} = x_0 x_1 \ldots$. We assume that given an initial state $x_0 \in X$, a control input sequence $\mathbf{u}^N = u_0 u_1 u_2 \ldots u_{N-1} \in U^N$ and a sequence of environment inputs $\mathbf{w}^N = w_0 w_1 w_2 \ldots w_{N-1} \in W^N$, the resulting

horizon-N run of a system modeled by equation (1), which we denote by

$$\xi(x_0, \mathbf{u}^N, \mathbf{w}^N) = (x_0 u_0 w_0)(x_1 u_1 w_1)(x_2 u_2 w_2)...(x_N u_N w_N),$$

is unique. In addition, we introduce a generic cost function $J(\xi(x_0, \mathbf{u}, \mathbf{w}))$ that maps (infinite and finite) runs to $\mathbb{R}$.

2.1 Signal Temporal Logic
We consider STL formulas defined recursively according to the grammar

$$\varphi ::= \pi^\mu \mid \neg\pi^\mu \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \square_{[a,b]} \ \psi \mid \varphi \ \mathcal{U}_{[a,b]} \ \psi$$

where π^μ is an atomic predicate $\mathbb{R}^n \to \mathbb{B}$ whose truth value is determined by the sign of a function $\mu : \mathbb{R}^n \to \mathbb{R}$ and ψ is an STL formula.

The validity of a formula φ with respect to the discrete-time signal $\mathbf{x}$ at time t_k, noted $(\mathbf{x}, t_k) \models \varphi$ is defined inductively as follows:

$$
\begin{aligned}
(\mathbf{x}, t_k) &\models \pi^\mu & &\Leftrightarrow & &\mu(x_k) > 0 \\
(\mathbf{x}, t_k) &\models \neg\pi^\mu & &\Leftrightarrow & &\neg((\mathbf{x}, t_k) \models \pi^\mu) \\
(\mathbf{x}, t_k) &\models \varphi \wedge \psi & &\Leftrightarrow & &(\mathbf{x}, t_k) \models \varphi \wedge (\mathbf{x}, t_k) \models \psi \\
(\mathbf{x}, t_k) &\models \varphi \vee \psi & &\Leftrightarrow & &(\mathbf{x}, t_k) \models \varphi \vee (\mathbf{x}, t_k) \models \psi \\
(\mathbf{x}, t_k) &\models \square_{[a,b]} \varphi & &\Leftrightarrow & &\forall t_{k'} \in [t_k+a, t_k+b], (\mathbf{x}, t_{k'}) \models \varphi \\
(\mathbf{x}, t_k) &\models \varphi \ \mathcal{U}_{[a,b]} \ \psi & &\Leftrightarrow & &\exists t_{k'} \in [t_k+a, t_k+b] \text{ s.t. } (\mathbf{x}, t_{k'}) \models \psi \\
& & & & &\wedge \forall t_{k''} \in [t_k, t_{k'}], (\mathbf{x}, t_{k''}) \models \varphi.
\end{aligned}
$$

A signal $\mathbf{x} = x_0 x_1 x_2 ...$ satisfies φ, denoted by $\mathbf{x} \models \varphi$, if $(\mathbf{x}, t_0) \models \varphi$. Informally, $\mathbf{x} \models \square_{[a,b]} \varphi$ if φ holds at every time step between a and b, and $\mathbf{x} \models \varphi \ \mathcal{U}_{[a,b]} \ \psi$ if φ holds at every time step before ψ holds, and ψ holds at some time step between a and b. Additionally, we define $\lozenge_{[a,b]} \varphi = \top \ \mathcal{U}_{[a,b]} \ \varphi$, so that $\mathbf{x} \models \lozenge_{[a,b]} \varphi$ if φ holds at some time step between a and b.

An STL formula φ is *bounded-time* if it contains no unbounded operators; the *bound* of φ is the maximum over the sums of all nested upper bounds on the temporal operators, and provides a conservative maximum trajectory length required to decide its satisfiability. For example, for $\square_{[0,10]} \lozenge_{[1,6]} \varphi$, a trajectory of length $N \geq 10+6 = 16$ is sufficient to determine whether the formula is satisfiable. This bound can be computed in time linear in the length of the formula.

2.2 Robust Satisfaction of STL formulas
Quantitative or robust semantics define a real-valued function ρ^φ of signal $\mathbf{x}$ and t such that $(\mathbf{x}, t) \models \varphi \equiv \rho^\varphi(\mathbf{x}, t) > 0$. In this work, we utilize a quantitative semantic for *space*-robustness, which is defined as follows:

$$
\begin{aligned}
\rho^{\pi^\mu}(\mathbf{x}, t_k) &= \mu(x_k) \\
\rho^{\neg\pi^\mu}(\mathbf{x}, t_k) &= -\mu(x_k) \\
\rho^{\varphi \wedge \psi}(\mathbf{x}, t_k) &= \min(\rho^\varphi(\mathbf{x}, t_k), \rho^\psi(\mathbf{x}, t_k)) \\
\rho^{\varphi \vee \psi}(\mathbf{x}, t_k) &= \max(\rho^\varphi(\mathbf{x}, t_k), \rho^\psi(\mathbf{x}, t_k)) \\
\rho^{\square_{[a,b]} \varphi}(\mathbf{x}, t_k) &= \min_{t_{k'} \in [t+a, t+b]} \rho^\varphi(\mathbf{x}, t_{k'}) \\
\rho^{\varphi \ \mathcal{U}_{[a,b]} \ \psi}(\mathbf{x}, t_k) &= \max_{t_{k'} \in [t+a, t+b]} (\min(\rho^\psi(\mathbf{x}, t_{k'}), \\
& \qquad\qquad \min_{t_{k''} \in [t_k, t_{k'}]} \rho^\varphi(\mathbf{x}, t_{k''}))
\end{aligned}
$$

To simplify notation, we denote ρ^{π^μ} by ρ^μ for the remainder of this paper. The robustness of satisfaction for an arbitrary STL formula is computed recursively from the above semantics in a straightforward manner, by propagating the values of the functions associated with each operand using min and max operators corresponding to the various STL operators. For example, the robust satisfaction of π^{μ_1} where $\mu_1(x) = x - 3 > 0$ at time 0 is $\rho^{\mu_1}(\mathbf{x}, 0) = x_0 - 3$. The robust satisfaction of $\mu_1 \wedge \mu_2$ is the minimum of ρ^{μ_1} and ρ^{μ_2}. Temporal operators are treated as conjunctions and disjunctions along the time axis: since we deal with discrete time, the robustness of satisfaction of $\varphi = \square_{[0,2]} \mu_1$ is

$$\rho^\varphi(x, t) = \min_{t_k \in [0,2]} \rho^{\mu_1}(x, t_k) = \min\{x_0 - 3, x_1 - 3, \ldots, x_K - 3\}$$

where $0 \leq t_0 < t_1 < \ldots < t_K \leq 2 < t_{K+1}$.

note that for continuous time, the min and max operations would be replaced by inf and sup, respectively.

The robustness score $\rho^\varphi(\mathbf{x}, t)$ should be interpreted as *how much* model $\mathbf{x}$ satisfies φ. Its absolute value can be viewed as the distance of $\mathbf{x}$ from the set of trajectories satisfying or violating φ, in the space of projections with respect to the functions μ that define the predicates of φ.

REMARK 1. *We have introduced and defined a Boolean and a quantitative semantics for STL over discrete-time signals, which can be seen as roughly equivalent to Bounded Linear Temporal Logic (BLTL). There are several advantages of using STL over BLTL. First, STL allows us to explicitly use real time in our specifications instead of integer indices, which we find more elegant. Second, although in the rest of this paper we will focus on the control of the discrete-time system Σ_d, our goal is to use the resulting controller for the control of the continuous system Σ. Hence the specifications should be independent from the sampling time Δt. Finally, note that the relationship between the continuous-time and discrete-time semantics of STL depending on discretization error and sampling time is beyond the scope of this paper. The interested reader can refer to [9] for further discussion on this topic.*

2.3 MILP Encoding for Controller Synthesis
In order to synthesize a run that satisfies a STL formula φ, we add STL constraints to a MILP formulation of the control synthesis problem as in [24]. We first represent the system trajectory as a finite sequence of states satisfying the model dynamics in equation (1). Then, we encode the formula φ with a set of MILP constraints; our encoding produces a MILP as long as the functions μ that define the predicates π^μ in φ are linear or affine.

2.3.1 Constraints on system evolution
The system constraints encode valid finite (horizon-N) trajectories for a system of the form (1) – these constraints hold if and only if the trajectory $\xi(x_0, \mathbf{u}_N, \mathbf{w}_N)$ satisfies (1). A typical situation is when f_d is linear. In that case, the constraints on system evolution are of the form

$$x_1 = Ax_0 + B_u u_0 + B_w w_0$$
$$x_2 = Ax_1 + B_u u_1 + B_w w_1$$
$$\dots$$
$$x_N = Ax_{N-1} + B_u u_{N-1} + B_w w_{N-1}$$

2.3.2 STL constraints

The robustness of satisfaction of the STL specification, as defined in Section 2.2, provides a natural objective for the MILP defined in Section 2.3, either in the absence of, or as a complement to domain-specific objectives on turns of the system.

As described in Section 2.2, the robustness of a STL specification φ can be computed recursively on the structure of the formula. Moreover, since max and min operations can be expressed in a MILP formulation using additional binary variables, this does not add complexity to the encoding (although the additional variables make it more computationally expensive in practice). For a given formula φ, we introduce a variable r_k^φ, and an associated set of MILP constraints such that $r_k^\varphi > 0$ if and only if φ holds at time t_k. Given φ, we recursively generate MILP constraints for every subformula or φ, such that r_0^φ determines whether φ holds in the initial state. For example, additionally, we enforce that the value of $r_k^\varphi = \rho^\varphi(\mathbf{x}, t_k)$. The reader is referred to [24] for details of this encoding. The advantage of this *robustness-based* encoding is that it allows us to maximize or minimize the value of r_0^φ, obtaining a trajectory that maximizes or minimizes the robustness of satisfaction.

The union of the STL constraints and system constraints yields a MILP, enabling us to check feasibility and find a solution when possible using a MILP solver; for further details and examples see [24]. Given an objective function on runs of the system, we can also find an optimal trajectory that satisfies the STL specification. The robustness provides a natural objective for this MILP, either in the absence of, or as a complement to domain-specific objectives on runs of the system.

Mixed integer-linear programs are NP-hard, and hence impractical when the dimensions of the problem grow. We present the computational costs of the above encoding in terms of the number of variables and constraints in the resulting MILP. If P is the set of predicates used in the formula and $|\varphi|$ is the length (i.e. the number of operators), then $O(N \cdot |P|) + O(N \cdot |\varphi|)$ continuous variables are introduced. In addition, $O(N)$ binary variables are introduced for every instance of a Boolean operator, i.e., $O(N \cdot |\varphi|)$ Boolean variables.

The dimensionality of the discrete-time system 1 affects the size of the constructed MILP linearly via the constraints encoding system evolution (more precisely, through the size of the set of predicates P). However, given the efficiency of modern MILP solvers, there is no evidence that a linear increase in the problem size would in practice lead to more than a linear increase in computational time for a solution. Methods based on abstraction or discretization of the state space, on the other hand, are much more likely to have exponential complexity with respect to system dimensionality.

We note, however, that our approach is more sensitive to the size of the specifications, and in particular to the nesting degree of temporal operators. We report on the scalability of our approach in Section 6.

3. PROBLEM STATEMENT

We address the problem of synthesizing control inputs for a system operating in the presence of potentially adversarial, *a priori* uncertain external inputs or disturbances. The controllers we produce will provide guarantees for specifications of the form $\varphi \doteq \varphi_e \Rightarrow \varphi_s$, where φ_e places assumptions on the external environment, and φ_s specifies desired guarantees on the plant behavior. In this work, φ_e refers exclusively to properties of signals $\mathbf{w} \in W^\omega$, whereas φ_s refers to properties of $\mathbf{x} \in \mathcal{X}^\omega$ and $\mathbf{u} \in U^\omega$.

We now formally state the synthesis problem for reactive controllers subject to STL specifications of the form above, and its receding horizon formulation.

PROBLEM 1 (STL REACTIVE SYNTHESIS). *Given a system of the form in equation (1), initial state x_0, trajectory length N, STL formula φ and cost function J, compute:*

$$\operatorname*{argmin}_{\mathbf{u}^N} \quad \max_{\mathbf{w}^N \in \{\mathbf{w} \in W^N | \mathbf{w} \models \varphi_e\}} \quad J(\xi(x_0, \mathbf{u}^N, \mathbf{w}^N))$$

$$s.t. \quad \forall \mathbf{w}^N \in W^N, \quad \xi(x_0, \mathbf{u}^N, \mathbf{w}^N) \models \varphi$$

PROBLEM 2 (RECEDING HORIZON REACTIVE SYNTHESIS). *Given a system of the form in equation (1), initial state x_0, STL formula φ and cost function J, at each time step k, compute:*

$$\operatorname*{argmin}_{\mathbf{u}^{H,k}} \quad \max_{\mathbf{w}^{H,k} \in \{\mathbf{w} \in W^H | \mathbf{w} \models \varphi_e\}} \quad J(\xi(x_k, \mathbf{u}^{H,k}, \mathbf{w}^{H,k}))$$

$$s.t. \quad \forall \mathbf{w} \in W^\omega, \quad \xi(x_0, \mathbf{u}, \mathbf{w}) \models \varphi,$$

where H is a finite horizon provided as a user input or selected in some other fashion, $\mathbf{u}^{H,k}$ is the horizon-H control input computed at each time step and $\mathbf{u} = u_0^{H,0} u_0^{H,1} u_0^{H,2}...$.

In Sections 4 and 5, we present both a finite-trajectory solution to Problem 1, and a solution to Problem 2 for a large class of STL formulas. A key component of our solution is to use our previously presented encoding of STL specifications as MILP constraints [24] in combination with MILP constraints representing the system dynamics to efficiently solve the resulting constrained optimization problem.

4. COUNTEREXAMPLE-GUIDED FINITE HORIZON SYNTHESIS

We propose a solution to Problem 1 using a counterexample guided inductive synthesis (CEGIS) procedure. We first consider bounded STL properties φ, bounded by $N \in \mathbb{N}$. Once we have this scheme for synthesizing control for finite trajectories satisfying bounded specifications, we will use a receding horizon scheme for infinite trajectories.

We now describe the steps of Algorithm 1 in detail. In Step 2, we choose an initial instance $\mathbf{w}^0$ of an environment that

---
Algorithm 1 CEGIS Algorithm for Problem 1

1: **procedure** CEGIS(ξ, x_0, N, φ, J)
2:     Let $\mathbf{w}^0 = (w_1^0, w_2^0, ... w_{N-1}^0)$, s.t. $\mathbf{w}^N \models \varphi_e$
3:     $W_{cand} = \{\mathbf{w}^0\}$
4:     **while** True **do**
5:

$$\mathbf{u}^0 \leftarrow \operatorname*{argmin}_{\mathbf{u} \in U^N} \quad \max_{\mathbf{w}^0 \in W_{cand}} (J(\xi(x_0, \mathbf{u}, \mathbf{w}^0)))$$
$$\text{s.t. } \forall \mathbf{w}^0 \in W_{cand}, \ \xi(x_0, \mathbf{u}, \mathbf{w}^0) \models \varphi_s,$$

6:     **if** $\mathbf{u}^0 ==$ null **then**
7:       Return INFEASIBLE
8:     **end if**
9:

$$\mathbf{w}^1 \leftarrow \operatorname*{argmin}_{\mathbf{w} \in W^N} \quad \rho^\varphi(\xi(x_0, \mathbf{u}^0, \mathbf{w}), 0)$$
$$\text{s.t. } \mathbf{w}^1 \models \varphi_e$$

10:     **if** $\rho^\varphi(\xi(x_0, \mathbf{u}^0, \mathbf{w}^1)) > 0$ **then**
11:       Return $\mathbf{u}^0$
12:     **else**
13:       $W_{cand} \leftarrow W_{cand} \cup \{\mathbf{w}^1\}$
14:     **end if**
15:     **end while**
16: **end procedure**
---

satisfies the specification φ_e. We do so using the open-loop synthesis algorithm for bounded-time STL described in [24]. Our initial set of candidate environment inputs is a singleton, $W_{cand} = \{\mathbf{w}^0\}$ (Step 3). Then, in Step 5, we compute the optimal control input $\mathbf{u}^0$ with respect to this environment, such that the system specification φ_s is satisfied; this step also uses the solution in [24]. If the problem in Step 5 is infeasible, we know that there is a control input $\mathbf{w}^0 \in W_{cand}$ against which no control input can satisfy φ, so we can stop and return (Step 7). Otherwise, in Step 9, we find an environment $\mathbf{w}^1$ that satisfies φ_e, but also minimizes the robustness of satisfaction of φ for the control input $\mathbf{u}^0$. Essentially, this step tries to find an environment that falsifies the specification φ when the control input $\mathbf{u}^0$ is used. If the minimum robustness $\rho^\varphi(\xi(x_0, \mathbf{u}^0, \mathbf{w}^1))$ thus computed is positive, this implies $\forall \mathbf{w} \in W^N \ \xi(x_0, \mathbf{u}^0, \mathbf{w}) \models \varphi$, so we can return the control input $\mathbf{u}^0$ as our result in Step 11. Otherwise, we have generated a counterexample to $\mathbf{u}^0$ being the desired control input, i.e. an environment $\mathbf{w}^1$ that falsifies φ when $\mathbf{u}^0$ is used. We use this counterexample to guide our inductive synthesis in Step 13, by adding it to the set of environments to be considered in the next iteration. We then resume execution of the while loop from Step 4.

THEOREM 1. *If Algorithm 1 returns $\mathbf{u}^N \in U^N$, then $\forall \mathbf{w}^N \in W^N$, $\xi(x_0, \mathbf{u}^N, \mathbf{w}^N) \models \varphi$. If Algorithm 1 returns INFEASIBLE, then Problem 1 is infeasible.*

Note that Algorithm 1 does not fully solve Problem 1, because it does not always ensure cost-optimality of $\mathbf{u}^N$ with respect to all disturbances $\mathbf{w}^N \in W^N$ — the returned $\mathbf{u}^N$ is optimal with respect to a specific set of disturbances $W_{cand} \subseteq W^N$.

Since $|W_{cand}|$ grows by 1 at every iteration of the while loop, the MILP in Step 5 grows linearly with the number of iterations, since we duplicate constraints for each new counterexample. If W is finite, W_{cand} will converge, and Algorithm 1 is sound and complete. Otherwise, we execute a maximum number of iterations of the while loop before declaring the problem infeasible.

In practice, solving the problem in Step 5 becomes expensive as W_{cand} grows, in particular because the objective is now non-linear. While state-of-the-art MILP solvers e.g. Gurobi[1] handle nonlinear objective functions efficiently, we can preserve the difficulty of the problem at each iteration by setting $W_{cand} = \{\mathbf{w}^1\}$ in Step 13 instead of growing the set of candidates. This breaks completeness even for finite sets W, since we may oscillate between two disturbances, but preserves soundness with respect to the satisfaction of φ, while allowing faster solutions at each iteration of the loop.

In the case studies described in Section 6, we find that a few number of iterations through the while loop suffices to either find a satisfying control input or render the problem infeasible.

5. RECEDING HORIZON SYNTHESIS
In this section, we will describe a solution to Problem 2 by adding STL constraints to a receding horizon control framework. At each step t of the computation, we will employ the CEGIS approach in Section 4 to find a finite trajectory of fixed horizon length H, such that the trajectory accumulated over time satisfies φ.

Note that this problem is relatively simple for bounded-time STL formulas φ, as described in [24]. Here the length of the horizon H is chosen to be at least the bound of formula φ. Then, at time step 0, we synthesize control $\mathbf{u}^{H,0}$ using the formulation in Section 4, and execute only the first time step $u_0^{H,0}$; we then observe $w_0^{H,0}$ and x_1. Then at the next step, we solve for $\mathbf{u}^{H,1}$, while constraining the values of $u_0^{H,1} = u_0^{H,0}, w_0^{H,1} = w_0^{H,0}$ in the MILP, and retaining the STL constraints on the trajectory up to time H. Keeping track of the history in this manner ensures that the formula is satisfied over the length-H prefix of the trajectory, while solving for $\mathbf{u}^{H,t}$ at every time step t.

Suppose that we have a specification $\psi = \Box \varphi$, where φ is a bounded-time formula with bound H. In this case, we can stitch together trajectories of length H using a receding horizon approach to produce an infinite computation that satisfies the STL formula. At each step of the receding horizon computation, we search for a finite trajectory of horizon length $2H$, keeping track of the past values and robustness constraints necessary to determine satisfaction of ψ at every time step in the trajectory.

First we define a procedure:

$$\text{CEGIS}^*(\xi, x_0, N, \psi = \Box \varphi, J, \mathbf{P}^H, \mathbf{u}_{old}^k)$$

that takes additional inputs $\mathbf{P} = \{P_0, P_1, ..., P_{H-1}\}$ and $\mathbf{u}_{old}^k = u_{old_0}^k u_{old_1}^k ... u_{old_{k-1}}^k$, and is identical to Algorithm 1, except that the optimization problem in Step 5 is solved

---
[1] http://www.gurobi.com/

with the added constraints:

$$\rho^\varphi(\xi(x_0, \mathbf{u}, \mathbf{w}^0), i) > P_i \quad \forall i \in [0, H-1]$$
$$\forall i < k, u_i = u^k_{old_i}$$

Algorithm 1 (CEGIS) solves reactive synthesis for bounded horizon formulas. We are designing Algorithm 2 to deal with unbounded formulas, by invoking bounded-horizon synthesis at each time step. To ensure soundness of this infinite-horizon synthesis algorithm, some history needs to be carried forth from one horizon to another to ensure consistency between the newly synthesized inputs and those produced in previous steps. This is achieved by the two additional arguments of CEGIS* and the corresponding added constraints. The first constraint enforces satisfaction of φ at all time steps $i \in [0, H-1]$. The second constraint fixes the first k values of the newly computed input to values computed in the previous time step.

Given CEGIS*, we define a receding horizon control procedure as in Algorithm 2. At each time step, we compute control inputs over a horizon of $2H$.

Algorithm 2 RHC Algorithm for Problem 2

1: **procedure** RHC($\xi, x_0, \psi = \square \varphi, J$)
2:     Let M be a large positive constant.
3:     Let H be the bound of φ.
4:     Set $P_0 = 0$ and $P_i = -M \quad \forall 0 < i \leq H$.
5:     Compute $\mathbf{u}^0 = u^0_0 u^0_1 .... u^0_{2H-1}$ as:

   $$\mathbf{u}^0 \leftarrow \text{CEGIS}^*(\xi, x_0, 2H, \square_{[0,H]} \varphi, J, \mathbf{P}^H, \emptyset)$$

6:     **for** $k = 1; k <= H; k = k+1$ **do**
7:         Set $\mathbf{u}^k_{old} = u^0_0 u^1_1 u^2_2 ... u^{k-1}_{k-1}$
8:         Set $P_i = 0$ for $0 \leq i \leq k, P_i = -M \quad \forall k < i \leq H$.
9:         Compute $\mathbf{u}^k = u^k_0 u^k_1 .... u^k_{2H-1}$ as:

   $$\mathbf{u}^k \leftarrow \text{CEGIS}^*(\xi, x_k, 2H, \square_{[0,H]} \varphi, J, \mathbf{P}^H, \mathbf{u}^k_{old})$$

10:     **end for**
11:     **while** True **do**
12:         Set $\mathbf{u}^k_{old} = u^{k-1}_1 u^{k-1}_2 u^{k-1}_3 ... u^{H-1}_H$.
13:         Set $P_i = 0$ for $0 \leq i \leq H$.

   $$\mathbf{u}^k \leftarrow \text{CEGIS}^*(\xi, x_k, 2H, \square_{[0,H]} \varphi, J, \mathbf{P}^H, \mathbf{u}^k_{old})$$

14:         $k = k+1$
15:     **end while**
16: **end procedure**

Algorithm 2 has two phases, a *transient* phase (Lines 4-10) and a *stationary* phase (Lines 11-14). The transient phase applies until an initial control sequence of length H has been computed, and the stationary phase follows. In the transient phase, the number of stored previous inputs ($\mathbf{u}^k_{old}$) as well as the number of time steps at which formula φ is enforced (i.e. time steps for which $P_i = 0$) grows by one at each iteration, until they both attain a maximum of H at iteration H. Every following iteration uses a window of size H for stored previous inputs, and sets all $P_i = 0$. The size-H window of previously-computed inputs advances forward one step in time at each iteration after step H. In this manner, we keep a record of the previously computed inputs required to ensure satisfaction of φ up to H time steps in the past.

THEOREM 2. *Let* $\mathbf{u}^*$ *be the infinite sequence of control inputs generated by setting* $u^*_k = u^k_H$, *where* $\mathbf{u}^k = u^k_0 u^k_1 ... u^k_{2H-1}$ *is the control input sequence of length $2H$ generated by Algorithm 2 at time t_k. Then* $\forall \mathbf{w} \in W^\omega, \xi(x_0, \mathbf{u}^*, \mathbf{w}) \models \psi$.

PROOF. Since H is the bound of φ, the satisfaction of φ at time k is established by the control inputs $u^*_k \ldots u^*_{k+H-1}$.

At time $k + H$,

$$\begin{aligned}
\mathbf{u}^{k+H}_{old} &= u^{k+H}_0 u^{k+H}_1 u^{k+H}_2 ... u^{k+H}_{H-1} \\
&= u^{k+H-1}_1 u^{k+H-1}_2 u^{k+H-1}_3 ... u^{k+H-1}_{H-1} u^{k+H-1}_H \\
&= u^{k+H-2}_2 u^{k+H-2}_2 u^{k+H-2}_4 ... u^{k+H-2}_H u^{k+H-1}_H \\
&= \cdots \\
&= u^k_H u^{k+1}_H u^{k+2}_H ... u^{k+H-1}_H
\end{aligned}$$

and so all the inputs required to determine satisfaction of φ at time t have been fixed. Moreover, if $\mathbf{u}^{k+H}$ is successfully computed, then by the correctness of Algorithm 1, $\mathbf{u}^{k+H}_{old}$ has the property that $\forall \mathbf{w}^H \in W^H, \xi(x_t, \mathbf{u}^{k+H}_{old}, \mathbf{w}^H) \models \varphi$. Since $u^k_H u^{k+1}_H u^{k+2}_H ... u^{k+H-1}_H = \mathbf{u}^{k+H}_{old}$, we see that $\forall \mathbf{w}^H \in W^H, \xi(x_k, u^*_k ... u^*_{k+h}, \mathbf{w}^H) \models \varphi$.

It follows that $\forall \mathbf{w}^\omega \in W^\omega, \xi(x_0, \mathbf{u}^*, \mathbf{w}) \models \varphi$. $\square$

We have therefore shown how a control input can be synthesized for infinite sequences satisfying ψ, by repeatedly synthesizing control for sequences of length $2H$. A similar approach applies for formulas $\Diamond \varphi$ and $\varphi \, \mathcal{U} \, \psi$, where φ, ψ are bounded-time.

6. CASE STUDIES
We now validate our approach in simulation, for case studies in building climate control and autonomous driving.

6.1 Building Climate Control
We consider the problem of controlling building indoor climate in a commercial building equipped with a HVAC system controlled by a receding horizon control scheme. We adopt the model proposed by Maasoumy et al. [20], and the receding horizon control formulation proposed by Maasoumy et al. [19], with the objective of minimizing the total energy cost (in dollar value).

As shown in Figure 1, we model a building with 4 rooms; we denote the temperature of room r_i by T_i, and that of the outside by T_5. The temperature of a room is governed by differential equations depending on properties such as the heat capacity, heat absorption, thermal resistance and area of the walls between the room and its neighboring rooms, the radiative heat flux density on external walls, heat capacity and air mass flow into the room, transmissivity of the glass of windows, the total area of the windows on walls surrounding the room and the internal heat generation in the room. Further details on this thermal model can be found in [20].

The heat transfer equations for each wall and room yield a system of the form $\dot{x} = f(x, u, w)$ where $x \in \mathbb{R}^n$ is the state vector representing the temperature of the nodes in the thermal network (including rooms and walls) and $u \in \mathbb{R}^{lm}$ is the input vector representing the air mass flow rate and discharge air temperature of conditioned air into each thermal

Figure 1: Building layout for HVAC control. We show results for the temperature in Room 1. Temperatures from the neighboring rooms and outside are treated as uncertain disturbances around nominal temperatures obtained from existing measurements.

zone (with l being the number of inputs to each thermal zone, e.g. air mass flow and supply air temperature). The vector w stores the estimated disturbance values, aggregating various unmodeled dynamics such as the outside temperature, internal heat generation and radiative heat flux density, and can be estimated using historical data [20]. In this work, we only show results for controlling the temperature of Room 1, and include the temperature of the neighboring rooms as part of these unmodeled dynamics w.

Assume that the system dynamics are discretized with a sampling time of Δt, and let H be the prediction horizon (in number of time steps). Here we consider $\Delta t = 0.5$ hr and $H = 12$. At each time t_k, the receding horizon controller solves an optimal control problem to compute $\mathbf{u}_k^H = u_k \dots u_{k+H-1}$, minimizing the cumulative norm of $\mathbf{u}_k^H$: $\sum_{k=0}^{H-1} \|u_k\|$. We assume a known occupancy function occ_k, which is equal to 1 when the room is occupied and to 0 otherwise. The purpose of the controller is to maintain a comfort temperature given by T^{comf} whenever the room is occupied, while minimizing the cost of heating. The assumption on the environment is that the disturbances w are in a range bounded by ϵ around some reference w^{ref}, obtained from historical data. Formally:

$$x_{k+1} = f_d(x_k, u_k, w_k)$$
$$\varphi_e = \square_{[0,H]}(|w_k - w_k^{\mathrm{ref}}| < \epsilon)$$
$$\varphi_s = \square_{[0,H]}((\mathrm{occ}_k > 0) \Rightarrow (T_k > T_k^{\mathrm{comf}}))$$
$$J(\xi(x_0, \mathbf{u}^H, \mathbf{w}^H)) = \sum_{k=0}^{H-1} \|u_k\|$$

The STL formula φ was encoded using the robust MILP encoding. Figure 2(b) presents results of executing the receding horizon controller synthesized using Algorithm 1, while modeling the disturbance as bounded (more precisely, satisfying φ_e) but non-deterministic; we used $\epsilon = 3$ (which corresponds to an uncertain variation of $3°F$ in the temperature of the neighboring rooms and outside). Compare this

Figure 3: Two vehicles crossing an intersection simultaneously. The red car is the *ego* vehicle, which we control. The black car is part of the adversarial environment.

with Figure 2(a), where the disturbance is modeled as corresponding exactly to w_t^{ref} (i.e. $\epsilon = 0$). In both cases, the actual disturbance (undepicted) was exactly w^{ref}.

We observe that the controller designed to operate in an adversarial environment is more conservative, and starts heating the room earlier (e.g. time step 4 instead of 5) in response to the same predicted occupancy signal, to account for the possibility of a higher disturbance. Additionally, when the occupancy signal is non-zero, the control input applied to counter the worst case disturbance results in a temperature that is higher than in the deterministic case; the result is that the temperature plot rises further above the minimum temperature of T_{comf} in the nondeterministic case.

As we previously observed in [24], most of the time is spent initially creating the MILP, while solving it takes a fraction of a second for each time step. In practice, the CEGIS loop of Algorithm 1 was executed fewer than 2 times for most time steps. As expected, the number of CEGIS iterations was greater for the case where $\epsilon = 3$ than $\epsilon = 0$, reflecting the greater nondeterminism.

The HVAC model used in this case study is 5-dimensional [20]; this represents a significant improvement over reactive synthesis techniques based on discrete abstraction, which do not typically scale past 2 or 3 continuous variables. We expect our techniques to scale well to higher dimensions. The main culprit when it comes to problem size is the length of the horizon required to ensure satisfiability. This increases with the nesting of temporal operators.

## 6.2 Autonomous Driving in Nondeterministic Environments

We now consider the problem of controlling an autonomous vehicle operating in the presence of other, potentially adversarial vehicles.

In this example, two moving vehicles approach an intersection, which they must cross. We let the red car in Figure 3 be the *ego* vehicle (the vehicle we control), and the black car be part of the environment. We define the state space

(a) Known Disturbances: $\epsilon = 0$

(b) Uncertain Disturbances: $\epsilon = 3$

Figure 2: Receding horizon control of Room 1 temperature under constraints based on occupancy, expressed in STL. In both plots, the black line in the topmost subplots represent the actual value of $T_{\mathbf{comf}}$, and the red line is the prediction of $T_{\mathbf{comf}}$ for the current planning horizon. The blue line in the plot of the occupancy signal represents the actual value, and the red dotted overlay is the prediction for the current horizon. The green dotted lines in the temperature and air flow plots depict the control input and resulting state trajectory computed during the current iteration of the receding horizon control computation. The blue lines in the control and temperature plots represent the portion of the control input that was actually executed and the resulting temperature T_1, respectively.

using a simplified 6-dimensional model, with the position of the two vehicles $((x^{\mathrm{ego}}, y^{\mathrm{ego}}), (x^{\mathrm{adv}}, y^{\mathrm{adv}}))$ and the velocity of the two $(v^{\mathrm{ego}} = \dot{y}^{\mathrm{ego}}, v^{\mathrm{adv}} = \dot{x}^{\mathrm{adv}})$ in ms^{-1} as state variables, and the acceleration $(a^{\mathrm{ego}} = \dot{v}^{\mathrm{ego}})$ of the ego vehicle as a single input. The disturbance is the acceleration of the environment vehicle $(a^{\mathrm{adv}} = \dot{v}^{\mathrm{adv}})$, which is allowed to take values in a bounded range. Thus:

$$\mathbf{x}_k = \begin{bmatrix} x_k^{\mathrm{ego}} & y_k^{\mathrm{ego}} & x_k^{\mathrm{adv}} & y_k^{\mathrm{adv}} & v_k^{\mathrm{ego}} & v_k^{\mathrm{adv}} \end{bmatrix}^{\top} \quad (2)$$

$$\mathbf{u} = a_k^{\mathrm{ego}} \quad \mathbf{w} = a_k^{\mathrm{adv}} \quad (3)$$

We assume each vehicle has the dynamics of a double integrator:

$$\begin{bmatrix} \dot{x}^{\mathrm{ego}} \\ \dot{y}^{\mathrm{ego}} \\ \dot{v}^{\mathrm{ego}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x^{\mathrm{ego}} \\ y^{\mathrm{ego}} \\ v^{\mathrm{ego}} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \mathbf{u} \quad (4)$$

$$\begin{bmatrix} \dot{x}^{\mathrm{adv}} \\ \dot{y}^{\mathrm{adv}} \\ \dot{v}^{\mathrm{adv}} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x^{\mathrm{adv}} \\ y^{\mathrm{adv}} \\ v^{\mathrm{adv}} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \mathbf{w} \quad (5)$$

Our specification in this example is that there should be no collisions at the intersection between the two vehicles, and that the ego vehicle's speed should be close to $1ms^{-1}$. Here the disturbance $\mathbf{w}$ is the acceleration of the adversary, whose

value is assumed to be close to a reference value, $\mathbf{w}^{\mathrm{ref}}$. We use the following STL formulas:

$$\varphi_e = \square(|\mathbf{w} - \mathbf{w}^{\mathrm{ref}}| < 0.1)$$
$$\varphi_s = \square(|y_k^{\mathrm{ego}} - x_k^{\mathrm{adv}}| < 2) \implies \square_{[0,2]}(|v_k^{\mathrm{ego}}| < 0.1)$$

The formula φ_s specifies that whenever y_k^{ego} is close to x_k^{adv}, i.e. within the range of $2m$, the ego vehicle should come to a stop $(|v_k^{\mathrm{ego}}| < 0.1)$ for a short period of time (2s). Figure 3, shows that the two vehicles will be close only when they are in the vicinity of the intersection. We expect the ego vehicle to stop at the intersection in order to allow the adversary to cross. In addition, we optimize the following cost function, which encourages the ego vehicle's speed to be close to $1ms^{-1}$.

$$J(\xi(x_0, \mathbf{u}^H, \mathbf{w}^H)) = \sum_{l=0}^{H-1} ||v_{k+l}^{\mathrm{ego}} - 1|| \quad (6)$$

Figure 4 illustrates the result of applying Algorithm 2 to synthesize control inputs for the ego vehicle. The first plot shows the position of the two vehicles, x_k^{adv} and y_k^{ego} (in m). The ego vehicle starts with a negative value on its y-axis $y_0^{\mathrm{ego}} < 0$, and the adversary starts with a positive x-value $x_0^{\mathrm{adv}} > 0$. Here the origin represents the middle of the intersection: at any time k if $y_k^{\mathrm{ego}} = x_k^{\mathrm{adv}} = 0$, the two cars have collided. The synthesized control input should therefore avoid such a collision, and the two vehicles should

Figure 4: Plot of position, velocity and acceleration of the ego and adversary vehicles. The second plot from the top shows the distance (in m) of each vehicle from the intersection. While the adversary vehicle drives straight through the intersection at a constant speed, the ego vehicle stops at t=8s, when it is around 2.5m from the intersection, then resumes moving around t=16s, thus avoiding collision.

not be at location 0 or its vicinity ($|y_k^{\mathrm{ego}} - x_k^{\mathrm{adv}}| < 2$) at the same time.

As seen in the first and second subplots in Figure 4, at time $t = 8s$, the ego vehicle stops at its current position in order to avoid collision with the adversary car. The vehicle proceeds after a short stop to let the adversary pass. The third subplot shows the velocity of the two vehicles, and the fourth plot represents the acceleration. Notice that the velocity of the ego vehicle stabilizes at $1ms^{-1}$ at most times as long as it avoids any collisions. The accelerations shown in the fourth plot include the control input synthesized using Algorithm 2, and the disturbance, i.e., the acceleration of the adversary.

7. DISCUSSION
The main contribution of this paper is a CEGIS procedure for synthesis of reactive controllers for systems satisfying STL specifications. We showed how our approach can be used as part of a receding horizon control scheme, to generate control for systems that must satisfy STL properties in the presence of adversarial environments, subject to domain-specific cost functions. We presented experimental results for controller synthesis on simplified models of a smart-building HVAC system and an autonomous car, and showed in simulation that the synthesized controllers satisfy the specified properties despite nondeterministic and adversarial environments.

8. ACKNOWLEDGEMENTS

9. REFERENCES
[1] H. Abbas, B. Hoxha, G. Fainekos, and K. Ueda. Robustness-guided temporal logic testing and verification for stochastic cyber-physical systems. In *Proc. of IEEE International Conference on CYBER Technology in Automation, Control, and Intelligent Systems*, 2014.

[2] R. Alur, R. Bodik, G. Juniwal, M. M. K. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *Proceedings of the IEEE International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, October 2013.

[3] E. Bartocci, L. Bortolussi, L. Nenzi, and G. Sanguinetti. On the robustness of temporal properties for stochastic models. In *Proceedings Second International Workshop on Hybrid Systems and Biology, HSB 2013, Taormina, Italy, 2nd September 2013.*, pages 3–19, 2013.

[4] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.

[5] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS*, pages 193–207, 1999.

[6] K. Chatterjee, T. A. Henzinger, R. Jhala, and R. Majumdar. Counterexample-guided planning. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence, Edinburgh, Scotland, July 26-29, 2005*, pages 104–111, 2005.

[7] X. C. Ding, M. Lazar, and C. Belta. LTL receding horizon control for finite deterministic systems. *Automatica*, 50(2):399–408, 2014.

[8] A. Donzé and O. Maler. Robust satisfaction of temporal logic over real-valued signals. In *FORMATS*, pages 92–106, 2010.

[9] G. Fainekos and G. Pappas. Robust sampling for MITL specifications. In J.-F. Raskin and P. Thiagarajan, editors, *Formal Modeling and Analysis of Timed Systems*, volume 4763 of *Lecture Notes in Computer Science*, pages 147–162. Springer Berlin Heidelberg, 2007.

[10] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas. Temporal logic motion planning for dynamic robots. *Automatica*, 45(2):343 – 352, 2009.

[11] G. E. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theor. Comput. Sci.*, 410(42):4262–4291, 2009.

[12] E. A. Gol and M. Lazar. Temporal logic model predictive control for discrete-time systems. In *Proceedings of the 16th international conference on Hybrid systems: computation and control, HSCC 2013, April 8-11, 2013, Philadelphia, PA, USA*, pages 343–352, 2013.

[13] X. Jin, A. Donzé, J. Deshmukh, and S. A. Seshia. Mining requirements from closed-loop control models. In *HSCC'13*, 2013.

[14] S. Karaman and E. Frazzoli. Vehicle routing problem with metric temporal logic specifications. In

Proceedings of the 47th IEEE Conference on Decision and Control, CDC 2008, December 9-11, 2008, Cancún, México, pages 3953–3958, 2008.

[15] S. Karaman and E. Frazzoli. Sampling-based motion planning with deterministic μ-calculus specifications. In *Proceedings of the 48th IEEE Conference on Decision and Control, CDC 2009, combined withe the 28th Chinese Control Conference, December 16-18, 2009, Shanghai, China*, pages 2222–2229, 2009.

[16] S. Karaman and E. Frazzoli. Linear temporal logic vehicle routing with applications to multi-UAV mission planning. *International Journal of Robust and Nonlinear Control*, 21(12):1372–1395, 2011.

[17] M. Kloetzer and C. Belta. A fully automated framework for control of linear systems from temporal logic specifications. *IEEE Transaction on Automatic Control*, 53(1):287–297, 2008.

[18] Y. Kwon and G. Agha. Ltlc: Linear temporal logic for control. In M. Egerstedt and B. Mishra, editors, *HSCC*, pages 316–329, 2008.

[19] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D. Callaway. Model predictive control approach to online computation of demand-side flexibility of commercial buildings hvac systems for supply following. In *IEEE American Control Conference (ACC 2014)*, Portland, USA, June 2014.

[20] M. Maasoumy Haghighi. *Controlling Energy-Efficient Buildings in the Context of Smart Grid: A Cyber Physical System Approach*. PhD thesis, University of California, Berkeley, Dec 2013.

[21] O. Maler and D. Nickovic. Monitoring temporal properties of continuous signals. In *FORMATS/FTRTFT*, pages 152–166, 2004.

[22] R. M. Murray, J. Hauser, A. Jadbabaie, M. B. Milam, N. Petit, W. B. Dunbar, and R. Franz. Online control customization via optimization-based control. In *In Software-Enabled Control: Information Technology for Dynamical Systems*, pages 149–174. Wiley-Interscience, 2002.

[23] P. Nuzzo, H. Xu, N. Ozay, J. Finn, A. Sangiovanni-Vincentelli, R. Murray, A. Donze, and S. Seshia. A contract-based methodology for aircraft electric power system design. *Access, IEEE*, PP(99):1–1, 2013.

[24] V. Raman, M. Maasoumy, A. Donzé, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia. Model predictive control with signal temporal logic specifications. In *Proc. of the IEEE Conf. on Decision and Control*, 2014.

[25] A. Solar-Lezama, L. Tancau, R. Bodík, S. A. Seshia, and V. A. Saraswat. Combinatorial sketching for finite programs. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 404–415. ACM Press, October 2006.

[26] E. M. Wolff, U. Topcu, and R. M. Murray. Optimization-based trajectory generation with linear temporal logic specifications. In *2014 IEEE International Conference on Robotics and Automation, ICRA 2014, Hong Kong, China, May 31 - June 7, 2014*, pages 5319–5325, 2014.

[27] T. Wongpiromsarn, U. Topcu, and R. M. Murray. Receding horizon temporal logic planning. *IEEE Trans. Automat. Contr.*, 57(11):2817–2830, 2012.