
Dynamics-Based Reactive Synthesis and Automated Revisions for High-Level Robot Control

Jonathan A. DeCastro ^{*†}

Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

Rüdiger Ehlers

Department of Computer Science, University of Bremen, 28359 Bremen, Germany

Matthias Rungger

Department of Electrical Engineering and Information Technology, Technical University of Munich, 80333 Munich, Germany

Ayça Balkan and Paulo Tabuada

Electrical Engineering Department, University of California, Los Angeles, Los Angeles, CA 90095, USA

Hadas Kress-Gazit

Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

Abstract

The aim of this work is to address issues where formal specifications cannot be realized on a given dynamical system subjected to a changing environment. Such failures occur whenever the dynamics of the system restrict the robot in such a way that the environment may prevent the robot from progressing safely to its goals. We provide a framework that automatically synthesizes revisions to such specifications that restrict the assumed behaviors of the environment and the behaviors of the system. We provide a means for explaining such modifications to the user in a concise, easy-to-understand manner. Integral to the framework is a new algorithm for synthesizing controllers for reactive specifications that include a discrete representation of the robot's dynamics. The new approach is demonstrated with a complex task implemented using a unicycle model.

* This work was supported by the NSF Expeditions in Computing project ExCAPE: Expeditions in Computer Augmented Program Engineering [grant number CCF-1138996].

† Corresponding author; e-mail: jad455@cornell.edu

1. Introduction

One of the benefits of the formal methods approach to robot mission planning is that it frees users from the burdens of programming controllers for complicated robot tasks, removing the arduous step of re-validation every time the task changes; see, e.g. Kloetzer & Belta (2008), Kress-Gazit et al. (2009), Tabuada & Pappas (2006), Wolff et al. (2013), Wongpiromsarn et al. (2010). Application of provably-correct controllers on fielded systems requires that the correctness guarantees extend to a possibly complex set of dynamics describing the robot. Recently, tools have been proposed for generating discrete abstractions for a wide class of nonlinear systems, e.g. Girard et al. (2010), Pola et al. (2008), Reißig (2011), Zamani et al. (2012). It is natural to exploit such abstractions when synthesizing controllers that are applicable to complex physical systems. This paper serves to solidify such a synthesis approach.

When composing tasks for applications such as self-driving cars, household robots, and disaster response robots, the user must specify both the desired system behaviors and any assumptions on the sensed environment. Often, such assumptions incorporate knowledge of the behavior of humans and other elements in the robot’s physical surroundings. If any of these assumptions conflict with the specified task in the context of a robot’s dynamics, the specification is considered to be *unrealizable*. That is to say, the environment is allowed to work against fulfilment of the specification by the robot assigned to the task. As the underlying cause for unrealizability is often difficult to determine, automated frameworks for debugging specifications (Könighofer et al. (2009), Raman & Kress-Gazit (2013)) and generating environment assumptions (Alur et al. (2013), Fainekos (2011), Li et al. (2011)) have gained attention recently.

This paper seeks to address the problem of realizability of specifications for robots with dynamics by introducing a novel framework that *automatically* suggests additions to the specification to resolve unrealizability issues. We do so by introducing a new synthesis approach that takes a specification agnostic to the dynamics of the robot, and synthesizes a controller (if possible) for a specific robot, given its discrete abstraction. If it turns out that the specification is unrealizable, we automatically generate revisions to the specification that render the task realizable with respect to the robot dynamics. These formulas effectively alter the environment assumptions and robot behaviors. The drawback of generating revisions independently of the user is that the resulting specifications may fail to capture the original intent of the mission. We address this by advocating for an means explaining such revisions in a manner that is simple to understand, permitting user interaction as the revisions are computed.

When considering dynamics, unrealizability can arise as a result of unwanted effects such as overshooting into undesired regions or an inability to make progress toward goals. Take the following scenario where rooms A and B are separated by a door:

Visit rooms A and B. Avoid the door if it is closed. Assume that, infinitely often, the door is open.

In the above scenario, the robot must be able to robustly avoid the walls of the workspace (boundaries) and appropriately react to the door as it opens and closes. If one applies this specification to robot with inertia, the task could fail because the robot may not be able to stop by the time the door closes, violating the statement “*Avoid the door if it is closed.*” By leveraging both the system dynamics and the specification, we could recover by adding the environment assumption “*The door must not close if within 1 meter of it.*” In this paper, we explore an automated method for suggesting such revisions.

The remainder of this paper is outlined as follows. In Section 2, we review relevant formal definitions and notation. We formally state the problem in Section 3, and address related work in the context of the problem in Section 4. In Section 5, we describe our approach for synthesis for specifications involving discrete abstractions. We present the approach for generating formulas and user feedback in Section 6. Next, we demonstrate our approach for a complex task carried out by an abstraction of a unicycle robot model in Section 7. Lastly, we provide a summary in Section 8.

2. Preliminaries

2.1. Linear Temporal Logic

Linear temporal logic (LTL) formulas are defined over the set AP of atomic propositions in the recursive grammar:

$$\varphi ::= \text{true} \mid \pi \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \bigcirc\varphi \mid \varphi_1 \text{U}\varphi_2$$

where π is an atomic proposition in AP . Respectively, \wedge and \neg are the Boolean operators “conjunction” and “negation”, and \bigcirc and U are the temporal operators “next” and “until”. From these operators, the following operators are derived: “disjunction” \vee , “implication” \Rightarrow , “equivalence” \Leftrightarrow , “always” \square , and “eventually” \diamond .

AP consists of a set of *environment* propositions \mathcal{X} describing the state of the discretized robot sensor values and a set of *system* propositions \mathcal{Y} describing the current pose of the robot in the workspace. For a set of propositions \mathcal{X} , let $\bigcirc\mathcal{X} = \{\bigcirc\pi_i\}_{\pi_i \in \mathcal{X}}$ be the set of propositions with the “next” operator. The LTL formulas are evaluated over infinite sequences $\sigma = \sigma_0\sigma_1\sigma_2\dots$ of truth assignments to the propositions in AP . σ is said to *satisfy* $\bigcirc\varphi$, $\square\varphi$, or $\diamond\varphi$ if φ holds true in the next position in the sequence, every position, or at some future position(s), respectively. We refer the reader to Vardi (1996) for a complete definition of the syntax and semantics of LTL formulas.

Definition 1 (Robot Mission Specification). The specifications we consider are LTL formulas of the form:

$$\varphi := \underbrace{(\varphi_i^e \wedge \varphi_t^e \wedge \varphi_g^e)}_{\varphi^e} \implies \underbrace{(\varphi_i^s \wedge \varphi_t^s \wedge \varphi_g^s)}_{\varphi^s}$$

The formulas φ_i^α , φ_t^α , and φ_g^α are defined over AP , where φ_i^α are formulas for the initial conditions, φ_t^α the allowed transitions (safety conditions) to be satisfied always, φ_g^α the goals (liveness conditions) to be satisfied infinitely often, and $\alpha = \{e, s\}$ (with e for ‘environment’ and s for ‘system’). The liveness guarantees take the form $\bigwedge_{i \in I_\alpha} \square \diamond (B_i^\alpha)$, where I_α is the index set of environment goals (B_i^e) or system goals (B_i^s).

Definition 2 (Controller Strategy and Execution). Define a *controller* as a finite-state machine $\mathcal{A} = (S, S_0, \mathcal{X}, \mathcal{Y}, \delta, \gamma_{\mathcal{X}}, \gamma_{\mathcal{Y}})$, where S is the set of controller states, $S_0 \subseteq S$ is the set of initial controller states, \mathcal{X} and \mathcal{Y} are sets of propositions described above, $\delta : S \times 2^{\mathcal{X}} \rightarrow S$ is a state transition relation providing the next state $s' \in S$ given the next state $s \in S$ and the current value of the environment input $z \in 2^{\mathcal{X}}$, i.e. $s' = \delta(s, z)$, $\gamma_{\mathcal{X}} : S \rightarrow 2^{\mathcal{X}}$ is a labelling function mapping controller states to the set of environment propositions evaluating to True for all transitions into that state, and $\gamma_{\mathcal{Y}} : S \rightarrow 2^{\mathcal{Y}}$ is a labelling function mapping controller states to the set of robot configuration propositions evaluating to True in that state.

Consider an infinite *execution* σ of \mathcal{A} , where $\sigma = (\gamma_{\mathcal{X}}(s_0), \gamma_{\mathcal{Y}}(s_0))(\gamma_{\mathcal{X}}(s_1), \gamma_{\mathcal{Y}}(s_1)), \dots$ for $s_0 \in S_0$ and $s_i \in S$, $i > 0$. φ is deemed *realizable* if there exists a finite-state machine \mathcal{A} such that every execution produced by \mathcal{A} satisfies φ . That is, at every $i \geq 0$, there exists an assignment of system variables \mathcal{Y} for all possible assignments of the environment variables \mathcal{X} such that σ satisfies φ . If there exist some environment behaviors on \mathcal{X} for which no such \mathcal{A} can be found, then φ is *unrealizable*.

A discrete controller can be synthesized from a specification follows by solving a two-player game played between the environment and the system, as described in Bloem et al. (2012).

2.2. Synthesis over a Connectivity Graph

A connectivity graph is an undirected graph describing those workspace regions that are accessible and adjacent to one another.

Given a bounded configuration space $W \subset \mathbb{R}^n$, let $\mathcal{R} = \{R_1 \dots R_p\}$ represent a set of disjoint regions whose closure covers W , where the open sets $R_i \subseteq W$. Let \mathcal{Y} be the set of propositions representing the workspace regions over which the topology model will be specified and $\pi_i \in \mathcal{Y}$ denote a region proposition that evaluates to True when the robot is in $R_i \in \mathcal{R}$.

Definition 3 (Topology Model). We define the *topology model* as a formula φ_t^{top} over \mathcal{Y} that encodes the allowed next regions given the current region. This formula is defined as follows:

$$\varphi_t^{top} = \bigwedge_{\pi_i \in \mathcal{Y}} \square \left(\pi_i \implies \bigvee_{\substack{\pi_j \in \mathcal{Y}: \\ cl(R_i) \cap cl(R_j) \neq \emptyset}} \bigcirc \pi_j \right),$$

where $cl(\cdot)$ denotes the closure operation on a set. Note that we also enforce mutual exclusion of regions; that is, the robot is only allowed to occupy one region at a time.

When combining the user specifications φ^e and φ^s with the topological model, we obtain the formula $\varphi^e \implies (\varphi^s \wedge \varphi_t^{top})$ written over $AP^{top} = \mathcal{X} \cup \mathcal{Y}$. A controller \mathcal{A} is synthesized if $\varphi^e \implies (\varphi^s \wedge \varphi_t^{top})$ is realizable.

2.3. Synthesis over Robot Abstractions

Our model of the behavior of the robot is a nonlinear differential equation

$$\dot{\xi}(t) = f(\xi(t), \nu(t)) \quad (1)$$

given by the function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, where $\xi(t)$ is the continuous state of the robot and $\nu(t)$ the command input of the robot at time $t \in \mathbb{R}_{\geq 0}$. We use $\xi_{x,\nu}$ to denote a *trajectory* of (1) with initial state x and input signal ν . We impose the usual regularity assumptions on f that imply the existence and uniqueness of solutions of (1). Moreover, we assume that (1) is forward-complete (its solution is defined for all $t \geq 0$), see e.g. Angeli & Sontag (1999).

We produce an estimate of the divergence of neighboring trajectories to obtain a conservative over-approximation of the continuous dynamics operating under a sample-and-hold controller framework. Let $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a continuous function, that satisfies $\beta(0, 0) = 0$ and $\beta(\cdot, t)$ is strictly increasing for every $t \in \mathbb{R}_{\geq 0}$. Let us assume that any pair of trajectories produced by the robot dynamics (1) satisfy

$$|\xi_{x,\nu}(t) - \xi_{x',\nu}(t)| \leq \beta(|x - x'|, t), \quad (2)$$

with $x, x' \in \mathbb{R}^n$, $t \in \mathbb{R}_{\geq 0}$ and constant input function $\nu \in U$. Given that (1) is Lipschitz continuous with constant L , the function $\beta(r, t) = re^{Lt}$ satisfies (2).

In constructing a discrete abstraction, we follow the approach in Pola et al. (2008), Reißig (2011), Zamani et al. (2012) for discretizing the bounded configuration space $W \subset \mathbb{R}^n$ and the bounded space of command inputs $U \subset \mathbb{R}^m$. We denote $[W]_\eta$ to be the uniform grid on W discretized with resolution η . This grid is defined as follows:

$$[W]_\eta := \{x \in W \mid \exists k \in \mathbb{Z}^n : x = k\eta\}.$$

We introduce the discretization parameters $\tau, \eta, \mu \in \mathbb{R}_{> 0}$, where η and μ define, respectively, the discretization of the robot configuration and command input spaces, and τ represents the sampling time.

We next denote a set of *configuration* propositions \mathcal{Y}_a describing the robot's pose, and a set of *locomotion command* propositions \mathcal{U}_a describing the controlled actions taken by the robot. We assume throughout that \mathcal{U}_a consists of purely robot

locomotion commands; the set can be generalized to capture other robot actions (e.g. wave hand, activate alarm), but this extension is omitted in this paper.

Definition 4 (Discrete Abstraction). The *discrete abstraction* $S_a(\tau, \eta, \mu)$ is defined by the tuple $(Q_a, V_a, \delta_a, \gamma_a^{\mathcal{Y}}, \gamma_a^{\mathcal{U}})$, where:

- $Q_a = [X]_\eta$ is the set of discretized robot configurations;
- $V_a = [U]_\mu$ is the set of discretized robot locomotion commands;
- $\delta_a : Q_a \times V_a \rightarrow 2^{Q_a}$ is a transition relation defined such that, there exists $q_a, q'_a \in Q_a$ and $v_a \in V_a, q'_a \in \delta_a(q_a, v_a)$ iff there exists $\xi_{x,\nu}$ with $\nu(t) = v_a$, for all $x \in \{x_0 \mid |x_0 - q_a| \leq \eta/2\}$ and $t \in [0, \tau)$, such that

$$|q'_a - \xi_{x,\nu}(\tau)| \leq \beta(\eta/2, \tau) + \eta/2; \quad (3)$$

- $\gamma_a^{\mathcal{Y}} : Q_a \rightarrow \mathcal{Y}_a$ labels each discretized configuration with the associated proposition in \mathcal{Y}_a that evaluates to True when the robot is in the given configuration;
- $\gamma_a^{\mathcal{U}} : V_a \rightarrow \mathcal{U}_a$ labels each discretized command with the associated proposition in \mathcal{U}_a that evaluates to True when the given command is active.

Additionally, we define $\gamma_a : \mathcal{R} \rightarrow 2^{\mathcal{Y}_a}$, a labelling function associating each region to a set of configuration propositions, as $\gamma_a(R_j) = \{\gamma_a^{\mathcal{Y}}(q_a) \mid q_a \in [R_j]_\eta\}$. Finally, we encode the transitions δ_a as a safety formula φ_t^a defined over $\mathcal{Y}_a \cup \bigcirc \mathcal{U}_a \cup \bigcirc \mathcal{Y}_a$ of the form

$$\varphi_t^a = \bigwedge_{\substack{\gamma_a^{\mathcal{Y}}(q_a) \in \mathcal{Y}_a \\ \gamma_a^{\mathcal{U}}(v_a) \in \mathcal{U}_a}} \square \left((\gamma_a^{\mathcal{Y}}(q_a) \wedge \bigcirc \gamma_a^{\mathcal{U}}(v_a)) \implies \bigvee_{\substack{\gamma_a^{\mathcal{Y}}(q'_a) \in \mathcal{Y}_a : \\ q'_a \in \delta_a(q_a, v_a)}} \bigcirc \gamma_a^{\mathcal{Y}}(q'_a) \right).$$

Suppose $\mathcal{Y}_a = \{\pi_0, \pi_1, \pi_2\}$, $\mathcal{U}_a = \{\pi_{fwd}\}$ and we start at grid cell π_0 at q_0 . Assume that the set of possible successor configurations under this command consists of cells 1 and 2; i.e. $\delta_a(q_0, v_0) = \{q_1, q_2\}$, where $\gamma_a^{\mathcal{Y}}(q_1) = \pi_1$, $\gamma_a^{\mathcal{Y}}(q_2) = \pi_2$ and $\gamma_a^{\mathcal{U}}(v_0) = \pi_{fwd}$. In this example, we write $\square((\pi_0 \wedge \bigcirc \pi_{fwd}) \implies \bigcirc(\pi_1 \vee \pi_2))$ as the corresponding safety formula.

2.4. Ensuring Correctness of the Continuous Behaviors

In the following, we briefly review the rules derived in Fainekos et al. (2009), Fainekos & Pappas (2009), Liu & Ozay (2014), Liu et al. (2012) for modifying LTL formulas φ^e , φ^s on AP^{top} to account for the behaviors of the continuous system in between sampling instants. Given the system (1), there exists a constant $b \in \mathbb{R}_{\geq 0}$ such that for all $\xi(t) \in W$, $u \in U$ and $t \in [0, \tau)$, we have

$$|f(\xi(t), u)| \leq b. \quad (4)$$

Through the process of inflation and deflation of physical workspace regions \mathcal{R} , we obtain the S_a -strengthened formulas $\hat{\varphi}^e$ and $\hat{\varphi}^s$. Conservative under-approximations \check{R}_i^ε (where regions are deflated or shrunk) and over-approximations \hat{R}_i^ε (where regions are inflated) may be constructed as follows:

$$\begin{aligned} \check{R}_i^\varepsilon &= \{x \in \mathbb{R}^n \mid \{x\} + \varepsilon\mathbb{B} \subseteq R_i\} \\ \hat{R}_i^\varepsilon &= \{x \in \mathbb{R}^n \mid x \in R_i + \varepsilon\mathbb{B}\}. \end{aligned}$$

Here we use $+$ to denote the Minkowski set addition and \mathbb{B} to denote the closed ball in \mathbb{R}^n with respect to the infinity norm $|\cdot|$ with radius τb . Supposing that the formulas φ^e , φ^s are in *negated normal form*, i.e., the negation only appears in front

of atomic propositions, we define the S_a -strengthened formulas $\hat{\varphi}^e$, $\hat{\varphi}^s$ by

$$\hat{\varphi}^\alpha = \varphi^\alpha[\neg\pi_{top,i}/\neg\gamma_a(\hat{R}_i^\varepsilon), \pi_{top,i}/\gamma_a(\hat{R}_i^\varepsilon)], \quad (5)$$

in which we are replacing the region propositions $\pi_{top,i} \in \mathcal{Y}$ with $\gamma_a(\hat{R}_i^\varepsilon)$ and $\neg\pi_{top,i}$ with $\neg\gamma_a(\hat{R}_i^\varepsilon)$.

By application of Theorem 1 of Liu & Ozay (2014) to the semantics in Fainekos et al. (2009), it is shown that, by choosing $\varepsilon \geq b\tau$, we are assured that the continuous trajectory of (1) satisfies $(\varphi^e \wedge \varphi_t^a) \implies \varphi^s$ if the sampled-time execution satisfies $(\hat{\varphi}^e \wedge \varphi_t^a) \implies \hat{\varphi}^s$.

Similar to Liu et al. (2013), we treat the possible non-determinism in the abstraction as adversarial: at each time-step, the environment assigns values to the environment \mathcal{X} . Then, the locomotion command is updated based on this input, and the abstraction chooses potential successor states to transition into given the current locomotion command. We therefore treat φ_t^a as a statement on the environment and adopt the formula $(\hat{\varphi}^e \wedge \varphi_t^a) \implies \hat{\varphi}^s$ over $AP = \mathcal{X} \cup \mathcal{Y}_a \cup \mathcal{U}_a$. We describe in Section 5 our procedure for checking realizability for this formula.

The synthesis of a discrete controller \mathcal{A}_{S_a} given the formula $\varphi^e \implies (\varphi^s \wedge \varphi_t^{top})$ and an abstraction S_a involves the following steps: (1) computation of the formula φ_t^a from S_a (Section 2.3); (2) computation of the S_a -strengthened formulas $\hat{\varphi}^e$, $\hat{\varphi}^s$ from the original formulas φ^e , φ^s (Section 2.4); and (3) extracting a strategy as is discussed in Section 5.

3. Problem Formulation

Let φ be a *realizable* formula, defined as

$$\varphi := \varphi^e \implies (\varphi^s \wedge \varphi_t^{top}),$$

in the set of atomic propositions AP^{top} under the assumption of a topological model φ_t^{top} , where the formulas φ^e and φ^s are, respectively, the user-defined environment assumptions and system guarantees. Additionally, we require that φ^e is not falsified by robot behaviors satisfying φ^s (i.e. no trivial behaviors). We refer to φ as a *general* formula. The goal is to synthesize controllers for such specifications.

Problem 1 (Synthesis under non-deterministic abstractions). Given the subformula φ_t^a for the robot abstraction and the S_a -strengthened formulas $\hat{\varphi}^e$ and $\hat{\varphi}^s$, synthesize a controller for the robot-specific specification

$$\varphi^{abs} := (\hat{\varphi}^e \wedge \varphi_t^a) \implies \hat{\varphi}^s.$$

in the set of propositions AP .

If φ^{abs} is *unrealizable*, determine a set of revisions to the specification that render the new specification realizable. A revision is a temporal logic formula that may be conjuncted with the original formula. Recall φ^{abs} is unrealizable if there exists some environment behavior admissible by the formula $\hat{\varphi}^e \wedge \varphi_t^a$ such that no system behaviors satisfy $\hat{\varphi}^s$.

Problem 2 (Generating revisions). Given a realizable φ , an unrealizable φ^{abs} , an abstraction S_a , and the S_a -strengthened formulas $\hat{\varphi}^e$ and $\hat{\varphi}^s$, automatically derive a set of formulas for the environment and the system such that

$$\varphi^{mod} := (\hat{\varphi}^e \wedge \varphi_t^a \wedge \psi_g^e \wedge \psi_t^e) \implies (\hat{\varphi}^s \wedge \psi_t^s) \quad (6)$$

is realizable and $\hat{\varphi}^e \wedge \varphi_t^a \wedge \psi_g^e \wedge \psi_t^e$ and $\hat{\varphi}^s \wedge \psi_t^s$ are satisfiable.

If such formulas exist, the user is provided with a suggested liveness assumption ψ_g^e and suggested set of safety assumptions and guarantees, ψ_t^e and ψ_t^s , respectively. For each transition formula, a simple-to-interpret statement summarizing each suggested revision is provided to the user. To illustrate the problem, consider the following example.

Example 1. Consider a robot tasked with fetching parts in a factory setting, illustrated in Figure 1. Here, *stockroom*, *station_1* and *station_2* are regions in the workspace belonging to \mathcal{Y} and *s1_occupied* and *s2_occupied* are sensors belonging to \mathcal{X} indicating when the respective region is occupied. The robot is required to visit the stockroom and the two workstations (system liveness) but avoid visiting those that are occupied (system safety). If the robot is within a workstation, the environment is required to keep that station unoccupied (environment safety). Also, the workstations are required to infinitely often be unoccupied (environment liveness).

The general specification φ is as follows:

$$\begin{aligned}
\Box \Diamond \textit{stockroom} \wedge \Box \Diamond \textit{station_1} \wedge \Box \Diamond \textit{station_2} & \triangleleft \text{sys liveness} \\
\Box \Diamond \neg \textit{s1_occupied} \wedge \Box \Diamond \neg \textit{s2_occupied} & \triangleleft \text{env liveness} \\
\Box (\textit{s1_occupied} \implies \bigcirc \neg \textit{station_1}) & \triangleleft \text{sys safety} \\
\Box (\textit{s2_occupied} \implies \bigcirc \neg \textit{station_2}) & \triangleleft \text{sys safety} \\
\Box (\textit{station_1} \implies \neg \textit{s1_occupied}) & \triangleleft \text{env safety} \\
\Box (\textit{station_2} \implies \neg \textit{s2_occupied}) & \triangleleft \text{env safety}
\end{aligned}$$

The specification φ is realizable. Suppose now we are given a fully-actuated planar robot governed by inertia, described by the system

$$\ddot{x} = u \quad \ddot{y} = v, \quad (7)$$

where $(u, v) \in U$ are robot commands and $(x, y) \in \mathbb{R}^2$ are the Cartesian robot coordinates. We derive an abstraction S_a of these dynamics in the configuration space $(x, y, \dot{x}, \dot{y}) \in W$ and obtain the formula φ^{abs} . With S_a , we want to synthesize a realizable controller for φ^{abs} that satisfies the task given an abstraction of these dynamics, with additional restrictions placed on the behaviors of the environment and the system.




Fig. 1. Resupply example.

Unrealizability can arise from a number of possible causes. One possibility is that the robot's inertia may prevent it from avoiding collisions with either one of the workstations upon sensing that it is occupied. Consider the case where *s1_occupied* is `False` and the robot is moving toward *station_1*. Suppose that, when moving, the robot must pass through two grid cells to decelerate to a stop. Then, the specification is unrealizable if *s1_occupied* is allowed to become `True`

when the robot is within two grid cells of *station_1*. This is an example of a *deadlock* behavior; the environment can force the system into certain states that have no legal transitions.

Another possibility is that the environment may toggle between states infinitely often, preventing the physical robot from making progress toward its goals. For instance, suppose the robot is approaching *station_1* or *station_2* and the environment toggles the values of *s1_occupied* and *s2_occupied* infinitely often. The robot in this case will be always changing directions, but unable to reach any of its goals before the environment changes once again. This is an example of a *livelock* behavior; the robot is prevented from reaching its goals as a result of repeating sequence of environment inputs. The behavior does not exist in the general formula because the topology graph always allows the robot to either remain in place or transit to an adjacent region once the environment has moved.

4. Related Work

We present a method that finds appropriate revisions to a specification when the inclusion of discrete abstractions cause unrealizability of specifications. Our work intersects with two lines of research: synthesis with dynamics, and assumption mining for reactive synthesis. We address related work in both of these areas below.

Synthesis with Dynamics. Tomlin et al. (2000) were among the first to generate verified controllers preserving safety and reachability specifications for nonlinear dynamical systems. In their approach, controllers are constructed by posing the problem as a differential game in hybrid systems (systems with mixed continuous and discrete states) analysis. The work of Mitchell et al. (2005) made such “reach-avoid” problems tractable through viscosity solutions to time-dependent partial differential equations.

Progress has been made more recently in developing techniques for synthesis of controllers satisfying specifications with more expressive requirements than reachability and safety. Namely, Kress-Gazit et al. (2009, 2011) introduce algorithms for reactive synthesis that take advantage of the *bisimulation property* allowing high-level controllers to be synthesized separately from the robot-specific low-level controllers. In order to bridge the gap between the high- and low-levels of abstraction, tools for automatically synthesizing controllers in the continuous domain based on high-level specifications have been introduced recently in DeCastro & Kress-Gazit (2014), Fainekos et al. (2006). Bhatia et al. (2010), Maly et al. (2013) have solved this problem using multi-layered synthesis approaches, where certain parts of the control strategy are left open for an online planner to complete at runtime.

Another perspective has been to incorporate the dynamics directly into the reactive synthesis process. Rather than having an additional step for synthesizing continuous controllers that adhere to the behaviors of high-level controller, Liu et al. (2013) introduce a discrete abstraction of the underlying dynamics (Girard et al. (2010), Reißig (2011), Zamani et al. (2012)) directly into the synthesis process. The problem of synthesizing controllers using abstractions of physical systems is well-studied: in provably-correct mission planning, researchers have considered robot dynamics ranging from simple single or double integrators (Fainekos et al. (2009), Kress-Gazit et al. (2009)) and piecewise linear models (Tumova et al. (2010), Yordanov et al. (2012)) to nonlinear (Bhatia et al. (2010), Girard et al. (2010), Wolff et al. (2013), Wongpiromsarn et al. (2010), Zamani et al. (2012)), switched (Liu et al. (2013)) and hybrid systems (Maly et al. (2013)). In Kloetzer & Belta (2008), non-deterministic abstractions are used to synthesize non-reactive LTL formulas using model checking methods. LTL synthesis for switched systems was considered in Liu & Ozay (2014), Liu et al. (2013), where the authors propose methods for computing fine-grained abstractions and switching protocol synthesis for reactive tasks. Our work is inspired by the nonlinear synthesis methods of Liu et al. (2013), Pola et al. (2008), Zamani et al. (2012). Where our approach differs is that our synthesis algorithm takes specifications written agnostic to the robot and produces controllers that are applicable to a given robot platform, given its discrete abstraction. We furthermore explore the case when the dynamics contribute to the unrealizability of such specifications and supply revisions in such cases.

Assumption Mining. Unrealizable specifications that are complicated to parse greatly benefit from automated tools for computing revisions that restrict the environment and system behaviors. Our approach for computing such revisions is closely related to recent methods described in Fainekos (2011), Li et al. (2011), and Alur et al. (2013). In Fainekos (2011), a method is devised for determining the cause of unrealizability for non-reactive tasks and providing specification recommendations to the user. In the reactive setting, Li et al. (2011) present a debugging method for unrealizable specifications based on templates (LTL formulas) mined from an environment counterstrategy. In essence, a counterstrategy captures the possible behaviors for the environment for which there are no safe system moves that allow it to fulfill its goals. The method in Alur et al. (2013) generates templates of the form “ \diamond ” and “ $\diamond \square$ ” automatically from the counterstrategy, yielding additional safety and liveness environment statements that remove all execution traces of the counterstrategy. The work of Li et al. (2014) apply the counterstrategy-based environment assumption mining technique to an early warning system in human-in-the-loop control systems, demonstrated in an autonomous driving scenario. By removing the behaviors present in the counterstrategy, the modified environment is restricted in such a way as to permit the system to realize its goals under the strengthened assumptions, but can sometimes lead to specifications that no longer match the user’s intent.

Finding specific portions of the counterstrategy that lead to unrealizability is in general difficult without additional information or without post-processing the counterstrategy. Raman & Kress-Gazit (2013) developed algorithms to extract cores (minimal LTL formulas) for explaining unrealizable specifications. Similar to their work, we aim to identify two types of counterstrategy behaviors: deadlock and livelock. Where the authors use an “unrolling depth” to discover counterstrategy states leading to livelock (a cycle of environment inputs locking the robot away from its goals), we apply a similar notion to finding states leading to deadlock (a counterstrategy state where no further system transitions exist).

Our approach differs from existing works in several other ways. First, we adopt different strategy than Alur et al. (2013), Li et al. (2011) for computing revisions, that depends on whether deadlock or livelock is being removed from the counterstrategy. Second, we devise a means for efficiently computing revisions that reduces the number of times a counterstrategy needs to be synthesized. Lastly, we explain the revisions in a simple-to-understand manner. For example, the user will be provided with feedback in a structured format: “The specification is realizable as long as a person is not present when the robot is within 2 meters of the Hallway.”

5. Reactive Synthesis Under Nondeterministic Robot Abstractions

We synthesize reactive controllers by solving a two player game carried out between the environment (player 1) and the system (player 2) using the fixed-point algorithm described in Bloem et al. (2012). As described in this section, we make modifications to this algorithm in order to accommodate the adversarial nature of the discrete abstraction and preserve a mapping between φ and φ^{abs} . We also distinguish the expressive properties of the proposed approach with alternative approaches (e.g. Liu et al. (2013)).

Definition 5 (Controller Strategies for Non-deterministic Discrete Abstractions). A controller in our abstractions-based paradigm is a non-deterministic finite-state machine $\mathcal{A}_{S_a} = (S, S_0, \mathcal{X}, \mathcal{Y}_a, \mathcal{U}_a, \delta, \gamma_{\mathcal{X}}, \gamma_{\mathcal{Y}}, \gamma_{\mathcal{U}})$, where $S, S_0, \mathcal{X}, \mathcal{Y}_a, \mathcal{U}_a$, and $\gamma_{\mathcal{X}}$ are as described in Definition 2, $\delta : S \times 2^{\mathcal{X}} \rightarrow 2^S$ is a state transition relation providing the set of possible states at the next position in the sequence given the current controller state and the current value of the environment input, $\gamma_{\mathcal{Y}} : S \rightarrow 2^{\mathcal{Y}_a}$ is a labelling function mapping controller states to a set of possible robot configurations evaluating to True for transitions into that state, and $\gamma_{\mathcal{U}} : S \rightarrow 2^{\mathcal{U}_a}$ is a labelling function mapping controller states to a set of possible commands in that state.

Note that each state captures the possibility of non-determinism in the abstraction, obviating the sets-of-sets definition for $\gamma_{\mathcal{Y}}$. The robot configuration may be regarded as being an additional environment input, but differs in that the value it takes is not defined until the system chooses a control input. We therefore define an execution as a sequence of moves made by the environment and system, where the environment gets an extra move once the system has moved; in particular

$\sigma = (\gamma_{\mathcal{X}}(s_0), \gamma_{\mathcal{U}}(s_0), \gamma_{\mathcal{Y}}(s_0))(\gamma_{\mathcal{X}}(s_1), \gamma_{\mathcal{U}}(s_1), \gamma_{\mathcal{Y}}(s_1)), \dots$, where the robot motion at the current position in the sequence is a result of the environment input and robot command at the same position.

The three-move formulation is in contrast to the two-move formulation of Liu et al. (2013), where the state of the dynamical system at the current position in the sequence depends on the environment input and command at the *previous* position. Recall that, as we require a controller \mathcal{A}_{S_a} given the general-purpose formula $\varphi^e \implies (\varphi^s \wedge \varphi_t^{top})$, the three-move formulation allows for there to be a mapping between the subformulas φ^e , φ^s and $\hat{\varphi}^e$, $\hat{\varphi}^s$. For example, the safety guarantee $\varphi_t^s = \Box((person \wedge \bigcirc \neg person) \implies \bigcirc \neg r_1)$ has no equivalent counterpart $\hat{\varphi}_t^s$ in the method of Liu et al. (2013). The reason is that both *person* and r_1 are environment variables, and the fact that the system is required to move *before* r_1 is decided forces r_1 to be delayed to the next position in the execution. The translation is therefore $\Box((person \wedge \bigcirc \neg person) \implies \bigcirc \bigcirc \neg r_1)$, which is not in GR(1). Our strategy, on the other hand, allows for the environment to take an additional move after the system has moved, thereby preserving the mapping without delaying choosing r_1 to the next step.

We are able to assert that φ^{abs} is realizable if, for all positions in the execution and for all possible valuations of the environment variables \mathcal{X} , there exists a command \mathcal{U}_a such that all executions of \mathcal{A}_{S_a} satisfy the specification for all \mathcal{Y}_a . φ^{abs} is unrealizable if there exists some environment behavior(s) such that no \mathcal{U}_a can be found for which all robot configurations yield executions that satisfy the specification.

In order to determine if φ^{abs} is realizable, our focus on the GR(1) (generalized reactivity (1)) fragment lends to efficient solutions, as described in Bloem et al. (2012). Solving the GR(1) game is carried out by first introducing a *game structure*, defined as follows.

Definition 6 (Game Structure). A *game structure* is represented by the tuple $\mathcal{G} = (\mathcal{V}, \theta, \rho_e, \rho_r, \rho_s, \varphi_{win})$, where:

- $\mathcal{V} = \mathcal{X} \cup \mathcal{U}_a \cup \mathcal{Y}_a$ is the set of proposition valuations representing the position in the game;
- $\theta \subseteq 2^{\mathcal{V}}$ is the set of initial positions;
- $\rho_e \subseteq 2^{\mathcal{V}} \times 2^{\mathcal{X}}$ is a transition relation defining the set of environment inputs allowed by φ_t^e at the next position given the proposition values at the current position in the game;
- $\rho_r \subseteq 2^{\mathcal{V}} \times 2^{\mathcal{U}_a} \times 2^{\mathcal{Y}_a}$ is a transition relation derived from the robot formula φ_t^a defining the set of allowed robot configurations at the next position given the command at the next position and the robot configuration at the current position in the game;
- $\rho_s \subseteq 2^{\mathcal{V}} \times 2^{\mathcal{V}}$ is a transition relation defining the set of commands and robot configurations allowed by φ_t^s at the next position given the proposition values at the current position in the game; and
- φ_{win} is the winning condition.

The algorithm in Bloem et al. (2012) proceeds by considering winning positions that satisfy the system and environment safety transition relations while leading the system toward its goals. With the robot abstraction, we are faced with the additional step of ensuring that the positions chosen are safe for all possible robot configurations. We therefore begin by specifying the following enforceable predecessor operator Pre:

$$\begin{aligned} \llbracket \text{Pre} \mathcal{W} \rrbracket &:= \{v \in 2^{\mathcal{V}} \mid \forall v_x \subseteq \mathcal{X}, \exists v_u \subseteq \mathcal{U}_a, \forall v_y \subseteq \mathcal{Y}_a : \\ &\quad ((v, v_x) \in \rho_e) \implies [((v, v_u, v_y) \in \rho_r) \implies \\ &\quad [((v, v_x, v_u, v_y) \in \rho_s) \wedge ((v_x, v_u, v_y) \in \llbracket \mathcal{W} \rrbracket)]]]\}. \end{aligned}$$

where $\llbracket \varphi \rrbracket$ denotes the set of positions for which the formula φ evaluates to True. Intuitively, the enforceable predecessor is a set of positions enforcing that, for all environment inputs satisfying the environment transition relation, there exists a robot command such that all robot configurations bound to the transitions of the abstraction yields behaviors that remain safe, as long as the successor positions are taken from the set \mathcal{W} .

We next define a set of winning positions \mathcal{V}_{win} based on the μ -calculus formula Bloem et al. (2012) $\mathcal{V}_{win} = \nu\mathcal{W}_1 \cdot \bigwedge_{i_s \in I_s} \mu\mathcal{W}_2 \cdot \bigvee_{i_e \in I_e} \nu\mathcal{W}_3 \cdot N_{i_s i_e}$, where $N_{i_s i_e} = (B_{i_s}^s \wedge \text{Pre}\mathcal{W}_1) \vee \text{Pre}\mathcal{W}_2 \vee (\neg B_{i_e}^e \wedge \text{Pre}\mathcal{W}_3)$ and ν and μ represent, respectively, the greatest and least fixpoint operators. This formula ensures that there is a move that places the system strictly closer to one of its goals or one in which the system prevents one of the environment goals. Note that the control strategy we derive must be consistent with the physical system. The strategy \mathcal{A}_{S_a} that we extract from the winning set of positions is therefore preventing from falsifying the discrete abstraction φ_t^a (for example, those in which the robot is commanded to move to a configuration beyond W) by choosing only those control actions that also satisfy φ_t^a .

5.1. Counterstrategies for Nondeterministic Robot Abstractions

When a specification is unrealizable, one may synthesize a counterstrategy to find a sequence of environment inputs and robot configurations that prevent the robot from fulfilling its specification.

Definition 7 (Environment Counterstrategies for Non-deterministic Discrete Abstractions). We define an *environment counterstrategy* as a finite-state machine $\mathcal{A}^c = (Q^c, Q_0^c, \mathcal{X}, \mathcal{Y}_a, \mathcal{U}_a, \delta^c, \gamma_{\mathcal{U}}^c, \gamma_{\mathcal{X}}^c, \gamma_{\mathcal{Y}}^c)$, where

- Q^c is the set of counterstrategy states;
- $Q_0^c \subseteq Q^c$ is the set of initial counterstrategy states;
- $\mathcal{X}, \mathcal{Y}_a$ and \mathcal{U}_a are sets of propositions in AP ;
- $\delta^c : Q^c \rightarrow 2^{Q^c}$ is a nondeterministic transition relation returning the set of counterstrategy states at the next position in the sequence given the current state;
- $\gamma_{\mathcal{U}}^c : Q^c \rightarrow 2^{\mathcal{U}_a}$ is a labelling function mapping counterstrategy states to the set of locomotion command propositions in \mathcal{U}_a evaluating to True for all transitions into that state;
- $\gamma_{\mathcal{X}}^c : Q^c \rightarrow 2^{\mathcal{X}}$ is a labelling function mapping counterstrategy states to the set of environment propositions in \mathcal{X} evaluating to True in that state, and;
- $\gamma_{\mathcal{Y}}^c : Q^c \rightarrow 2^{\mathcal{Y}_a}$ is a labelling function mapping counterstrategy states to the set of robot configurations in \mathcal{Y}_a evaluating to True in that state.

We furthermore define $\delta^{c^{-1}} : 2^{Q^c} \rightarrow Q^c$ as the inverse transition relation mapping counterstrategy states to a set of predecessors; i.e. $\delta^{c^{-1}}(q') = \{q \in Q^c \mid q' \in \delta^c(q)\}$.

Our approach to synthesizing counterstrategies resembles that of Könighofer et al. (2009), where we use a fixed point computation to determine realizability of the specification, then extract a strategy \mathcal{A}^c from the set of positions that are winning for player 1. We define the enforceable predecessor operator for the counterstrategy Pre^c as follows:

$$\begin{aligned} \llbracket \text{Pre}^c \mathcal{W} \rrbracket &:= \{v \in 2^{\mathcal{V}^c} \mid \exists v_x \subseteq \mathcal{X}, \forall v_u \subseteq \mathcal{U}_a, \exists v_y \subseteq \mathcal{Y}_a : \\ &\quad ((v, v_x) \in \rho_e) \wedge [((v, v_u, v_y) \in \rho_r) \wedge \\ &\quad ((v, v_x, v_u, v_y) \in \rho_s)] \implies ((v_x, v_u, v_y) \in \llbracket \mathcal{W} \rrbracket)\}. \end{aligned}$$

The set of winning positions $\mathcal{V}_{win}^c = \mathcal{X} \cup \mathcal{U}_a \cup \mathcal{Y}_a$ for player 1 may then be computed from the formula $\mathcal{V}_{win}^c = \mu\mathcal{W}_1 \cdot \bigvee_{i_s \in I_s} \nu\mathcal{W}_2 \cdot \bigwedge_{i_e \in I_e} \mu\mathcal{W}_3 \cdot N_{i_s i_e}^c$, where

$$N_{i_s i_e}^c = (\neg B_{i_s}^s \vee \text{Pre}^c \mathcal{W}_1) \wedge \text{Pre}^c \mathcal{W}_2 \wedge (B_{i_e}^e \vee \text{Pre}^c \mathcal{W}_3),$$

A counterstrategy is synthesized by storing the sets $N_{i_s i_e}^c$ for each i_s and i_e at the last pass of the fixed point operation. Starting at the initial conditions φ_i^e, φ_i^s , for each state we identify an index i_s of the liveness guarantee that is currently being prevented by the counterstrategy. At a particular counterstrategy state q , we determine a successor q' similar to Könighofer et al. (2009) as follows. First, given the position v at state q , we fix an assignment of inputs $v'_x \in \mathcal{X}$. We next determine,

for this fixed assignment, the set of winning configurations $v'_y \in \mathcal{Y}_a$ that belong to $N_{i_s i_e}^c$ for every $v'_u \in \mathcal{U}_a$. A check is made to determine if a liveness assumption has been fulfilled at q ; if it has, then a new liveness assumption i_e is selected. For notational convenience, we define the set of successor values winning for player 1 at a counterstrategy state q as $M_{\mathcal{X}}(q) = \{v'_x \in \mathcal{X} \mid \forall v'_u \in \mathcal{U}_a, \exists v'_y \in \mathcal{Y}_a : (v, v'_x, v'_u, v'_y) \in N_{i_s(q) i_e(q)}^c, v = \gamma_{\mathcal{X}}^c(q) \cup \gamma_{\mathcal{U}}^c(q) \cup \gamma_{\mathcal{Y}}^c(q)\}$. This set contains the set of inputs for which there exist command-configuration combinations that are winning for player 1.

6. Generating Revisions to Unrealizable Specifications

In this section, we formalize our solution strategy for Problem 2. In similar fashion to Alur et al. (2013), Li et al. (2011), we make use of environment counterstrategies in order to search for environment assumptions that render the specification realizable. Our approach is outlined as follows: (1) from the specification φ^{abs} , we synthesize counterstrategies; (2) for each counterstrategy found, we compute environment and system transition subformulas ψ_t^e and ψ_t^s that prevent transitions to states in the counterstrategy from which the robot has no safe transitions (deadlock); (3) if a counterstrategy is found that is free of deadlock states, we compute liveness assumptions ψ_g^e that restrict transitions to cycles of states preventing the robot from fulfilling its goals (livelock). We introduce the following example to illustrate the major concepts discussed in this section.

Example 2. Consider the workspace shown in Figure 2a. Given $\mathcal{X} = \{sen\}$ where sen is the sensor input and $\mathcal{Y} = \{r1, r2\}$, we write a specification φ requiring the robot to visit $r2$ (lower-left gray region) when s is `False`, but avoid $r2$ when sen is `True`. Formally:

$$\begin{array}{ll}
 \square \diamond r2 & \triangleleft \varphi_g^s \\
 \square \diamond \neg sen & \triangleleft \varphi_g^e \\
 \square (\bigcirc sen \implies \bigcirc \neg r2) & \triangleleft \varphi_t^s \\
 \square (r2 \implies \bigcirc \neg sen) & \triangleleft \varphi_t^e \\
 \text{True} & \triangleleft \varphi_i^s \\
 \text{True} & \triangleleft \varphi_i^e
 \end{array}$$

The controller satisfying this specification is given in Figure 2b.

Given an abstraction where $\mathcal{Y}_a = \{x1, \dots, x16\}$ is an encoding of the set of 2-D robot configurations, and $\mathcal{U}_a = \{N, S, E, W\}$ are the robot commands for motion in the four cardinal directions, we derive a new specification φ^{abs} , where:

$$\begin{array}{ll}
 \square \diamond (x9 \vee x13) & \triangleleft \hat{\varphi}_g^s \\
 \square \diamond \neg sen & \triangleleft \hat{\varphi}_g^e \\
 \square (\bigcirc sen \implies \bigcirc \neg (x9 \vee x13)) & \triangleleft \hat{\varphi}_t^s \\
 \square ((x9 \vee x13) \implies \bigcirc \neg sen) & \triangleleft \hat{\varphi}_t^e \\
 \text{True} & \triangleleft \hat{\varphi}_i^s \\
 \text{True} & \triangleleft \hat{\varphi}_i^e
 \end{array}$$

The abstraction φ_t^a appears as arrows in the figure.




Fig. 2. 2-D example. (a) shows the workspace map and grid whose cells are labeled with the configuration variable. The white grid cells denote r_1 , while the gray denote r_2 . (b) shows the synthesized controller for φ .

6.1. Adding Revisions for Preventing Deadlock

To prevent deadlock, we introduce a scheme to process a counterstrategy and extract a set of environment assumptions that remove deadlock behaviors. Consider a counterstrategy \mathcal{A}^c whose deadlock states are collected in Q_{dead} . There exists an execution that eventually reaches a deadlock state q_{dead}^i ; specifically when $\exists q' \in \delta^c(q_{dead}^i) : q' = \emptyset$. We formally state this behavior as $\bigvee_{q^j \in \delta^{c^{-1}}(q_{dead}^i)} \diamond (\psi_1(q^j) \wedge \bigcirc \psi_2(q_{dead}^i))$, where $\psi_1(q)$ and $\psi_2(q)$ are propositional representations for the subsets of positions at counterstrategy state q :

$$\begin{aligned} \psi_1(q) &= \bigwedge_{\pi \in \gamma_{\mathcal{U}}^c(q) \cup \gamma_{\mathcal{Y}}^c(q)} \pi \wedge \bigwedge_{\pi \in (\mathcal{U}_a \cup \mathcal{Y}_a) \setminus (\gamma_{\mathcal{U}}^c(q) \cup \gamma_{\mathcal{Y}}^c(q))} \neg \pi, \\ \psi_2(q) &= \bigwedge_{\pi \in \gamma_{\mathcal{X}}^c(q)} \pi \wedge \bigwedge_{\pi \in \mathcal{X} \setminus \gamma_{\mathcal{X}}^c(q)} \neg \pi. \end{aligned}$$

In words, $\psi_1(q)$ captures the command and configuration moves at state q , and $\psi_2(q)$ captures the environment input at q .

To remove the environment behaviors in the counterstrategy causing deadlock, we adopt the following formula:

$$\bigwedge_{q^j \in \delta^{c^{-1}}(q_{dead}^i)} \square (\psi_1(q^j) \implies \bigcirc \neg \psi_2(q_{dead}^i)). \quad (8)$$

Before conjuncting each computed formula with ψ_t^e , a check is made to determine if it falsifies the left-hand side of φ^{mod} , i.e. there is no transition that satisfies $\hat{\varphi}^e \wedge \varphi_t^a \wedge \psi_t^e$. If this is the case, the formula is discarded and it is not included as a conjunct in ψ_t^e .

Example 3. Returning to Example 2, suppose we obtain a counterstrategy containing the states q_0, q_1, q_2, q_3, q_4 , as pictured in Figures 3a and 3b, starting in cell x_3 with $sen = \text{False}$. One of the possible executions in this counterstrategy eventually

leads the robot to cell $x4$ with the sensor $sen = \text{True}$:

$$\sigma = (\{\emptyset\}, \{\emptyset\}, \{x3\}), (\{\emptyset\}, \{W\}, \{x2\}), (\{\emptyset\}, \{W\}, \{x1\}), \\ (\{\emptyset\}, \{S\}, \{x5\}), (\{s\}, \{\emptyset\}, \{\emptyset\}).$$

where the i th time step corresponds to q_i . In this execution, the sensor sen remains **False** until the robot enters $x5$, at which point a transition in φ_t^s is violated. Hence q_4 is a deadlock state. The formula $\Diamond(\bigcirc sen \wedge \neg x1 \wedge S)^1$ is extracted by evaluating $\Diamond(\psi_1(q_4) \wedge \bigcirc \psi_2(q_4))$. The complement of this formula, $\Box((\neg x1 \wedge S) \implies \bigcirc \neg sen)$, is added as an additional environment assumption. This assumption negates the behaviors in the counterstrategy for that particular deadlock state. Being that there is only one deadlock state, we add no further assumptions. Upon adding this revision to the environment assumptions, we determine that the modified specification is realizable.

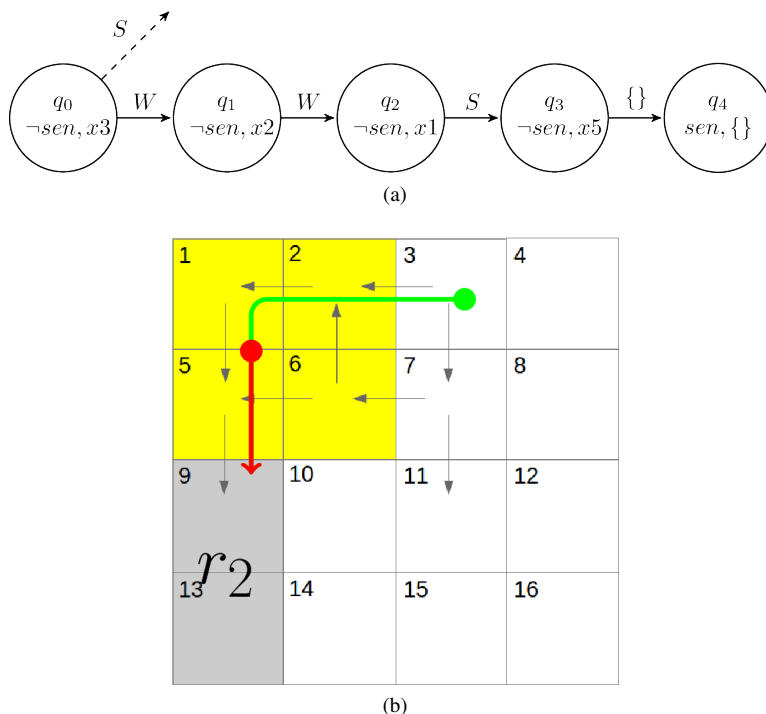


Fig. 3. (a) shows a partial counterstrategy for Example 3 leading to deadlock; (b) shows a corresponding robot trajectory leading to deadlock, where the green part of the path denotes where $sen = \text{False}$ and the red denotes where $sen = \text{True}$. The cells shaded yellow indicate configurations in which there are no sequence of commands that avoid reaching $r2$ eventually.

Notice that the controller synthesized in Example 3 produces executions that satisfy the specification but the system now assumes that the environment will always turn sen **False** whenever it reaches $x5$. In the example, consider the behavior when the robot starts at $x7$ with sen **True**. The execution of the robot in this case is

$$\sigma = (\{sen\}, \{\emptyset\}, \{x7\}), (\{sen\}, \{W\}, \{x6\}), (\{sen\}, \{W\}, \{x5\}), \\ (\{sen\}, \{S\}, \{x9\}), (\{\emptyset\}, \{E\}, \{x10\}), (\{\emptyset\}, \{N\}, \{x6\}), \dots$$

In this execution, sen remains **True** and, as the robot moves toward $r2$, the environment eventually must set sen to **False** to be consistent with the added assumption. When outside of $x5$, the robot follows the same sequence of moves regardless

¹ We only make the **True** action explicit (S in this case), since mutual exclusion disallows the other actions from being activated at the same time.

of the environment. Note that the controller for the original, realizable specification φ (Figure 2b) does not exhibit this behavior because there is no imposition on how the environment must behave based on the robot's configuration. In that case, if *sen* is True, the robot waits in *r1* until *sen* becomes False.

We remove such behaviors by disallowing a system transition whenever the same conditions found for the deadlock state hold at the state previous to deadlock. The idea is to force the system to react conservatively to the newly-added environment revision as if it were a deadlock state. Thus, when the robot is at a configuration previous to the deadlock state and the deadlock conditions hold currently, it will be forbidden from entering the configuration prior to deadlock. On the other hand, when the environment is not currently producing the same conditions as at deadlock, ψ_t^e forces the environment to “play fair” with the robot by not producing those conditions in the next step once the robot has made its move.

Formally, we disallow the behavior $\bigvee_{q^j \in \delta^{c-1}(q_{dead}^i)} \diamond (\bigcirc \neg \psi_2(q_{dead}^i) \wedge \bigcirc \neg \psi_1(q^j))$ by introducing an additional revision on φ_t^s

$$\bigwedge_{q^j \in \delta^{c-1}(q_{dead}^i)} \square (\bigcirc \neg \psi_2(q_{dead}^i) \implies \bigcirc \psi_1(q^j)). \quad (9)$$

Such a revision places a safety restriction on the robot, preventing it from entering a *neighboring* state to a deadlock state whenever the environment is set to the same value for which deadlock occurs. Doing this produces a specification that makes the system's behavior conservative; we are strengthening the conditions under which the robot may enter the neighboring state, when in fact the robot is not in any true danger of violating the original safety guarantees in φ until it reaches *r2*. Nonetheless, if the specification is realizable, the system will be able to react to the environment as long as the actions/configurations are not included in those specified in $\bigwedge_{q^j \in \delta^{c-1}(q_{dead}^i)} \psi_1(q^j)$.

If the modified formula is determined to be unrealizable and new deadlock states are found at a state $q^j \in \delta^{c-1}(q_{dead}^i)$, then we once again return to the original set of circumstances specified in Problem 2. We repeat the process in this section for as many times as required to eliminate deadlock states or when the specification is unrealizable. We may, however, avoid repeated synthesis of counterstrategies by applying the assumption and guarantee revisions explained above to entire *subtraces* of a single counterstrategy (a finite word of an execution trace for the counterstrategy). To do this, we identify states for which there is no safe command to be taken such that there exists a subtrace that eventually visits states in $Q^c \setminus Q_{dead}$. The search for deadlock revisions then reduces to a graph search on the counterstrategy, as summarized in Algorithm 1. The algorithm builds up a set of *deadlock-committed* states Q_{commit} by adding, via a breadth-first search (BFS in line 4), predecessor counterstrategy states from deadlock Q_{dead} for which all locomotion commands lead to states in Q_{commit} . For generating revisions and providing user feedback, we also maintain a mapping $Q_{reach} : Q_{commit} \rightarrow 2^{Q_{dead}}$ of deadlock states reachable from each $q \in Q_{commit}$. The search continues until a fixed point of states is reached where no additional deadlock-committed states can be found, at which point BFS returns a tuple containing Q_{commit} and Q_{reach} . The precise condition under which the search terminates is when a $q \in Q^c$ is found such that:

$$\exists q' \in \delta^c(q) : q' \notin Q_{commit}.$$

Here, Q_{commit} plays the role of Q_{dead} . We therefore replace Q_{dead} in the safety revisions (8) and (9) with Q_{commit} . To be precise, we replace (8) with

$$\bigwedge_{q^j \in Q_{commit}^i} \square \left(\psi_1(q^j) \implies \bigvee_{q^k \in Q_{reach}(q_{commit}^i)} \bigcirc \neg \psi_2(q^k) \right) \quad (10)$$

and (9) with

$$\bigwedge_{q^j \in Q_{commit}^i} \square \left(\bigwedge_{q^k \in Q_{reach}(q^i)} (\bigcirc \psi_2(q^k) \implies \bigcirc \neg \psi_1(q^j)) \right), \quad (11)$$

for each $q^i_{commit} \in Q_{commit}$. Consider the example below.

Algorithm 1 Computing deadlock-committed states.

```

procedure commitStates( $Q_{dead}$ )
  Initialize  $Q_{new}, Q_{commit}, Q_{reach}$  to  $Q_{dead}$ 
  while  $Q_{new} \neq \emptyset$  do
    ( $Q_{new}, Q_{reach}$ )  $\leftarrow$  BFS( $\mathcal{A}_c, Q_{commit}, Q_{reach}$ )
5:    $Q_{commit} \leftarrow Q_{commit} \cup Q_{new}$ 
  end while
  return  $Q_{commit}, Q_{reach}$ 
end procedure

```

Example 4. Starting from the result of Example 3, we compute a set of four deadlock-commit states $Q_{commit} = \{q_1, q_2, q_3, q_4\}$ corresponding to the cells $\{x6, x2, x1, x5\}$. We obtain the following ψ_t^e formulas:

$$\square((x5 \wedge S) \implies \bigcirc \neg sen) \quad (12)$$

$$\square((x1 \wedge S) \implies \bigcirc \neg sen) \quad (13)$$

$$\square((x2 \wedge W) \implies \bigcirc \neg sen) \quad (14)$$

$$\square((x6 \wedge N) \implies \bigcirc \neg sen), \quad (15)$$

and the following ψ_t^s formulas:

$$\square(\bigcirc sen \implies \bigcirc \neg(x5 \wedge S)) \quad (16)$$

$$\square(\bigcirc sen \implies \bigcirc \neg(x1 \wedge S)) \quad (17)$$

$$\square(\bigcirc sen \implies \bigcirc \neg(x2 \wedge W)) \quad (18)$$

$$\square(\bigcirc sen \implies \bigcirc \neg(x6 \wedge N)). \quad (19)$$

With these revisions added to ψ_t^e and ψ_t^s (highlighted orange in Figure 4, the modified specification eliminates the deadlock states present in the original counterstrategy. Additionally, note that none of the revisions falsify the environment.

Upon synthesis, we find that a counterstrategy synthesized from this modified specification does not contain deadlock states. In the next section, we discuss an approach to render the specification realizable by eliminating livelock behaviors.

6.2. Adding revisions for Preventing Livelock

In cases where the environment satisfies its own liveness conditions, the environment may be allowed to play in such a way that as the system cycles through an infinite sequence of states, the environment always keeps it away from one of its goals. Consider the behavior of the robot when the above ψ_t^e and ψ_t^s formulas (12)–(19) are introduced as revisions. Starting at

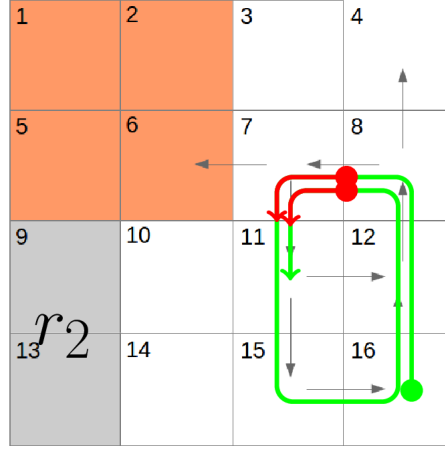


Fig. 4. Map showing configurations for which the revisions ψ_i^e and ψ_i^s from Example 4 apply; a counterstrategy execution trace, as explained in Section 6.2. The green part of the path denotes where $sen = \text{False}$ and the red denotes where $sen = \text{True}$.

x_{16} , the following behavior is possible:

$$\begin{aligned} \sigma = & (\{\emptyset\}, \{\emptyset\}, \{x_{16}\}), (\{\emptyset\}, \{N\}, \{x_{12}\}), (\{\emptyset\}, \{N\}, \{x_8\}), \\ & (\{sen\}, \{W\}, \{x_7\}), (\{\emptyset\}, \{S\}, \{x_{11}\}), (\{\emptyset\}, \{S\}, \{x_{15}\}), \\ & (\{\emptyset\}, \{E\}, \{x_{16}\}), (\{\emptyset\}, \{N\}, \{x_{12}\}), \dots \end{aligned}$$

In this execution (shown in Figure 4), the robot eventually cycles indefinitely between six cells in the workspace. Whenever the robot visits the cell x_7 , the environment activates sen , forcing the robot to move S to avoid violating the safety guarantee revision in (19). The environment is then able to satisfy its liveness goal ($\square \diamond (\neg sen)$), while preventing the robot from achieving its goal of reaching r_2 .

Once we obtain a counterstrategy free of deadlock states, we generate environment assumptions that remove the counterstrategy executions that exhibit livelock. The idea is to selectively choose states in the counterstrategy for which the robot still has winning actions to take and then apply liveness assumptions at those states to ensure that, always eventually, the robot is allowed to take these actions. We employ the result of the innermost fixed point computation $M_{\mathcal{X}}(q)$ stored when solving the counterstrategy game to find any states for which there exists an assignment of environment inputs that are not winning for the environment. Note that, for these inputs, there exists a command the robot may take which is winning for the system. We then prevent the environment from always making such assignments at these counterstrategy states.

Formally, for all $q \in Q^c$, our goal is to find a subset $Q_{cut} \subseteq Q^c$ for which $v'_x \notin M_{\mathcal{X}}(q)$, where $M_{\mathcal{X}}(q)$ are the set of environment inputs at a state $q \in Q^c$ that are winning for player 1. Consequently, for any $q \in Q_{cut}$ there are some $v'_x \subseteq \mathcal{X}$ and $v'_u \subseteq \mathcal{U}_a$ that lead to a position that is not winning for player 1 (i.e. $v'_x \notin M_{\mathcal{X}}(q)$). One can think of Q_{cut} as being those counterstrategy states where it is possible that the environment has been able to “cut away” a command that will allow the robot to proceed to its next goal by applying some environment input.

Using Q_{cut} , we formulate a set of liveness assumptions that restrict the environment from *always* behaving in a manner that prevents the system’s progress toward its goals. Notice that Q_{cut} contains all states for which there is an environment and system move not in player 1’s strategy; however, not all such moves are necessarily winning for player 2. For instance, a state in Q_{cut} could yield an environment input that does not allow player 1 to move strictly closer to its goal yet only allow system moves that place the system further away from its goal.

We therefore form a set $P_{cut} \subseteq Q_{cut}$ for which the robot has safe commands that are winning for the system. We use Q_{commit} (from the deadlock counterstrategy) to define the set of states where there exist system moves that lead the system

closer to its goals. We populate P_{cut} as follows:

$$P_{cut} = \{q \in Q_{cut} \mid \exists v_a \in V_a, \exists q' \in Q_{commit}, \exists q_a \in \delta_a(v_a, \gamma_y^c(q)) : \forall \pi \in \mathcal{Y}_a, \pi \in \gamma_a^y(q_a) \text{ iff } \pi \in \gamma_y^c(q')\}. \quad (20)$$

We then apply the environment liveness assumption

$$\Box \Diamond \bigvee_{q^i \in P_{cut}} \left(\psi_1(q^i) \wedge \bigwedge_{q^j \in \delta^c(q^i, \gamma_y^c(q^i))} \bigcirc \neg \psi_2(q^j) \right). \quad (21)$$

This liveness formula disallows the environment from always behaving in a way that denies the system from taking action that lead it closer to its goals, when the robot is in a configuration where there is such an action to be taken.

Example 5. With the specification φ^{abs} in Example 2 along with the revision (12)–(19), a counterstrategy is extracted as pictured in Figure 6. The set Q_{cut} consists of four states, $\{q_3, q_5, q_6, q_8\}$. Of these, $\{q_3, q_6, q_8\}$ are states in which the robot has an action (move W) taking the robot closer to r_2 that the environment can prevent. These are found from Example 3 and are collected in P_{cut} . We also find state $\{q_5\}$, for which there is an environment move keeping it from immediately realizing an environment goal ($\Box \Diamond (\neg sen)$) but does not lead the system closer to its goal. We apply environment liveness revisions ψ_g^e to the set P_{cut} :

$$\Box \Diamond ((x7 \wedge W \wedge \bigcirc \neg sen) \vee (x3 \wedge W \wedge \bigcirc \neg sen) \vee (x7 \wedge S \wedge \bigcirc \neg sen)). \quad (22)$$

The three above formulas correspond to regions appearing as blue regions in Figure 5. Adding this final revision produces a specification φ^{mod} that is realizable.




Fig. 5. Map showing regions associated with cut states from Example 5.

6.3. User Feedback

Before modifying the specification with the computed deadlock revisions, we alert the user of the consequences of these revisions. Given this information, the user may choose to accept these if they are, in fact, consistent with the original design intent. We provide feedback in the form of statements such as “Keep sensor `sen` `False` if the robot enters to within N meters of r_2 ” instead of raw LTL formulas.

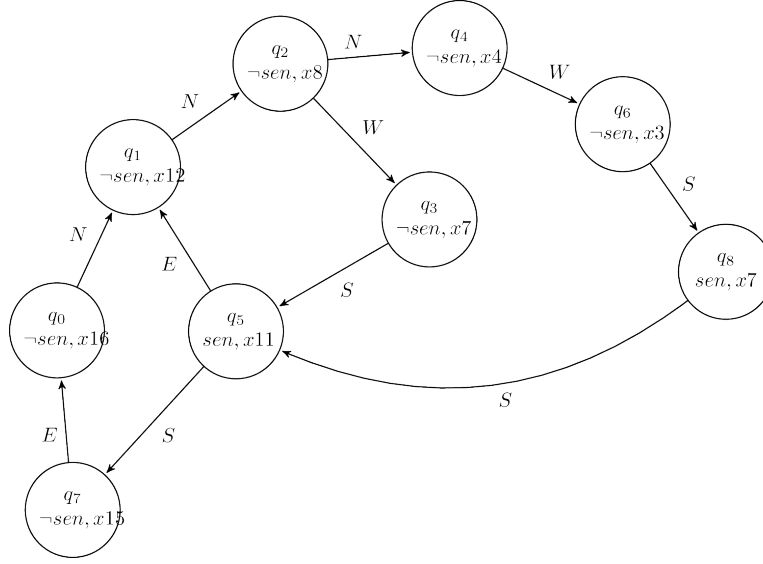


Fig. 6. Deadlock-free counterstrategy for Example 5.

For the combined environment and system revisions for deadlock, we use the mapping between cells and regions from the robot abstraction to assist in forming this metric. For each labeled workspace region $R_i \in \mathcal{R}$, we mark those deadlock states (if any exist) from whose predecessors there exists a robot command $v_a \in V_a$ (Definition 4) to reach R_i . Those marked deadlock states are collected in the set $P_{dead}(R_i)$, defined formally as:

$$P_{dead}(R_i) = \{q \in Q_{dead} \mid \forall q' \in \delta^{c^{-1}}(q), \exists v_a \in V_a, \exists q'_a \in \delta_a(v_a, \gamma_y^c(q')) : \\ \forall \pi \in \mathcal{Y}_a, \pi \in \gamma_a^y(q'_a) \text{ iff } \pi \in \gamma_a(R_i)\}.$$

In the fixed point computation in Algorithm 1, we keep track of each deadlock state reachable from each state added to Q_{commit} . We use this stored information to find the distances between each robot configuration in Q_{commit} . Rather than provide the user with a detailed set of conditions under which the environment would be required to adhere for the generated revisions to be satisfied, we advocate for simplicity by computing a conservative upper-bound on robot configurations where the environment restrictions must hold. For each R_i corresponding to a robot configuration that satisfies some deadlock state in Q_{dead} , we find the relative proximity to a deadlock condition (in terms of physical coordinates) by finding the maximal pairwise distance between any states affected by the deadlock revisions:

$$(q_i^*, q_{dead,i}^*) = \arg \max_{\substack{q \in Q_{commit}, \\ q' \in Q_{reach}(q) \cap P_{dead}(R_i)}} |\gamma_y^c(q) - \gamma_y^c(\delta^{c^{-1}}(q'))|. \quad (23)$$

Here, $|v_y|$ is the Euclidean norm of the real-valued abstraction state $q_a \in Q_a$ represented by a set of propositions $v_y \subseteq \mathcal{Y}_a$ that are True in that state. The pair of counterstrategy states q_i^* and $q_{dead,i}^*$ are those corresponding to a revision for region R_i where the distance is greatest, under the constraint that $q_{dead,i}^*$ is a deadlock state that is reachable from $q_i^* \in Q_{commit}$. Note that the distance between the configurations of the two states is:

$$dist_i = |\gamma_y^c(q_i^*) - \gamma_y^c(\delta^{c^{-1}}(q_{dead,i}^*))|.$$

The final step is to correlate each unique region R_i to the environment proposition assignments prevented by the safety assumption revisions ψ_i^e . Those prevented assignments are given in the formula $\psi_2(q_{dead,i}^*)$. That is, the added

environment assumptions prevent the environment from triggering the combination $\psi_2(q_{dead,i}^*)$. The data provided to the user is represented by the triple $(R_i, dist_i, \psi_2(q_{dead,i}^*))$. The triple can be displayed to the user as follows: “If the robot is within $dist_i$ of workspace region R_i , then the generated deadlock revisions (for a given counterstrategy) will be satisfied if the environment is not set to $\psi_2(q_{dead,i}^*)$.” Note that this metric supplies a sufficient but not necessary condition for satisfying the revisions. That is, there might be executions where the system enters within $dist_i$ with any environment setting yet still be able to satisfy the revision formulas.

Example 6. In the result of Example 4, let q_{dead} be the deadlock state computed by the counterstrategy corresponding to the configuration x_9 , and designate $Q_{commit} = \{q_1, q_2, q_3, q_4\}$ as the set of commit states for this deadlock. For workspace region r_2 , $P_{dead}(r_2) = q_{dead}$, and $Q_{reach}(q_i) = q_{dead}$ for $i = 1, \dots, 4$. We next determine the pair $(q_2^*, q_{dead,2}^*)$ to be

$$(q_2^*, q_{dead,2}^*) = \arg \max \left\{ \left| \begin{pmatrix} 0 \\ 1 \end{pmatrix} - \begin{pmatrix} 0 \\ 2 \end{pmatrix} \right|, \left| \begin{pmatrix} 0 \\ 0 \end{pmatrix} - \begin{pmatrix} 0 \\ 2 \end{pmatrix} \right|, \left| \begin{pmatrix} 1 \\ 0 \end{pmatrix} - \begin{pmatrix} 0 \\ 2 \end{pmatrix} \right|, \left| \begin{pmatrix} 1 \\ 1 \end{pmatrix} - \begin{pmatrix} 0 \\ 2 \end{pmatrix} \right| \right\} = (q_2, q_{dead}),$$

where the subscript 2 in the $*$ variables is used to signify the fact that the variables apply to region r_2 . Assuming $\eta = 1m$, the corresponding distance is $dist_2 = \sqrt{1^2 + 2^2} = 2.2m$. Finally, reflective of the revisions in (12)–(19), we note the subformula $\psi_2(q_{dead,2}^*) = sen$.

Therefore, the LTL formulas in (12)–(19) are summarized as: “If the robot enters to within $2.2m$ of r_2 , never set environment variable sen to True.”

6.4. Summary of the Approach

The approach described in Sections 5.1–6.3 may be combined into an automatic approach for synthesis assisted by user feedback. We outline this process in Algorithm 2. The approach has three main stages: computing revisions for deadlocks, composing an explanation of the revisions, and computing revisions for livelocks. Realizability is checked at every iteration of a while loop, terminating either when a specification is realizable or there are no further revisions to be computed (in this case, the counterstrategies from one iteration to the next will be identical). Revisions for eliminating deadlock in the counterstrategy are computed by applying (10) and (11) to the deadlock-committed states Q_{commit} . If these revisions falsify the environment and system, they are removed. The second step is to provide feedback to the user. Depending on the user’s response, the revisions are either applied or discarded.

The third step in the approach is to generate revisions for liveness as computed in (21). Once a candidate liveness assumption is computed, it is checked in Lines 32–39 to ensure that the system’s strategy does not contain a sequence of moves that cause the new liveness condition to be falsified. In such cases, `realizable` returns `False`, and the candidate liveness is removed. We refer the reader to Algorithm 5 of Raman & Kress-Gazit (2013) for further details. The user may elect to accept or discard this formula. Note that, if the specification is unrealizable and the counterstrategy is the same between iterations of the while loop, this means that no revisions have been found that meet the user’s criteria or do not falsify the specification. In this case, the algorithm terminates with an *unrealizable* output.

7. Example

In this section, we demonstrate the revision synthesis approach in an example scenario in which revisions to an unrealizable task specification are computed and added to the specification based on guidance provided by the user. We generate abstractions using the Pessoa Toolbox.² For synthesis, we use the Slugs Synthesis Tool, part of the LTLMoP Toolkit;³ the code used in this paper may be found at <https://github.com/jdc1177/slugs>.

² <https://sites.google.com/a/cyphylab.ee.ucla.edu/pessoa/>

³ <https://github.com/LTLMoP/slugs>

Algorithm 2 Synthesizing revisions for an unrealizable specification φ^{abs}

```

procedure synthRevisions( $\varphi^{abs}$ )
  Initialize  $\psi_t^e, \psi_t^s, \psi_g^e$  to True
  ( $realiz, \mathcal{A}_c^1, M_{\mathcal{X}}$ )  $\leftarrow$  ctrStrategy( $\varphi^{abs}$ )
   $\mathcal{A}_c^0 \leftarrow$  False,  $m \leftarrow 0$ 
5:  while  $\neg realiz \wedge (\mathcal{A}_c^m \neq \mathcal{A}_c^{m-1})$  do
       $Q_{commit} \leftarrow$  commitStates( $\mathcal{A}_c$ )
      for all  $q_{commit}^i \in Q_{commit}$  do
           $\psi_{t,cand}^e, \psi_t^e \leftarrow$  Eq. (10)
10:       $\psi_{t,cand}^s, \psi_t^s \leftarrow$  Eq. (11)
          if  $\neg(\varphi^e \wedge \varphi_t^a \wedge \psi_g^e \wedge \psi_t^e \wedge \psi_{t,cand}^e)$  or  $\neg(\hat{\varphi}^s \wedge \psi_t^s \wedge \psi_{t,cand}^s)$  then
               $\psi_t^e \leftarrow \psi_t^e \setminus \psi_{t,cand}^e$ 
               $\psi_t^s \leftarrow \psi_t^s \setminus \psi_{t,cand}^s$ 
          end if
15:      end for
           $\triangleright$  Step 1: Eliminate deadlocks

          for all  $R_i \in \mathcal{R}$  do
              ( $q_i^*, q_{dead}^*$ )  $\leftarrow$  Eq. (23)
               $dist_i \leftarrow |\gamma_{\mathcal{Y}}^c(q^*) - \gamma_{\mathcal{Y}}^c(q_{dead}^*)|$ 
20:              print ( $R_i, dist_i, \psi_2(q_{dead,i}^*)$ )
          end for
          if user accepts deadlock revisions then
               $\varphi^{mod} \leftarrow$  Eq. (6)
              ( $realiz, \mathcal{A}_c^m, M_{\mathcal{X}}$ )  $\leftarrow$  ctrStrategy( $\varphi^{mod}$ )
25:          end if
           $\triangleright$  Step 2: User feedback

           $Q_{cut} \leftarrow \{q \in Q^c \mid \exists v'_x \notin M_{\mathcal{X}}(q)\}$ 
           $P_{cut} \leftarrow$  Eq. (20)
           $\psi_{g,cand}^e, \psi_g^e \leftarrow$  Eq. (21)
30:          if  $\neg(\varphi^e \wedge \varphi_t^a \wedge \psi_g^e \wedge \psi_t^e \wedge \psi_{g,cand}^e)$  then
               $\psi_g^e \leftarrow \psi_g^e \setminus \psi_{g,cand}^e$ 
          else
               $\varphi^{try} \leftarrow$  Eq. (6)
               $realiz \leftarrow$  realizable( $\varphi^{try}$ )
35:          if  $\neg realiz$  then
               $\psi_g^e \leftarrow \psi_g^e \setminus \psi_{g,cand}^e$ 
          end if
           $\triangleright$  Step 3: Eliminate livelocks

          end if
          end if
40:          if user accepts livelock revisions then
               $\varphi^{mod} \leftarrow$  Eq. (6)
              ( $realiz, \mathcal{A}_c^m, M_{\mathcal{X}}$ )  $\leftarrow$  ctrStrategy( $\varphi^{mod}$ )
          end if
           $m++$ 
45:      end while
      return  $\varphi^{mod}$ 
end procedure

```

We return to the factory scenario in Example 1 using the workspace in Figure 1. To carry out this task, we select a robot described by a unicycle model that is governed by the kinematic relationship:

$$\dot{x} = v \cos \theta, \quad \dot{y} = v \sin \theta, \quad \dot{\theta} = \omega,$$

where the x and y are the Cartesian displacements in meters, θ is the orientation angle, and v and ω are, respectively, the forward and angular velocity inputs to the system. The car model is subjected to the constraint where it may only move with *positive* forward velocity (it cannot stop). An abstraction is generated for the three-dimensional configuration space and two-dimensional input space consisting of 2.2×10^6 states, with the chosen values $\eta = 0.15$, $\mu = 0.2$, $\tau = 0.35$.

The general specification is realizable, producing the controller pictured in Figure 7; however the specification φ^{abs} (with respect to the unicycle model) is unrealizable. With the approach in Algorithm 2, we compose revisions that render φ^{mod} realizable. After a counterstrategy is synthesized, revisions are found for a total of 2040 states in the counterstrategy (taking 1020 seconds to synthesize on a laptop PC with a dual-core processor and 8GB memory). A metric for these revisions is generated and the user is prompted with the following:

```
Deadlock revisions found.
When within 1.32 m of station_1, never set environment variable s1_occupied to True.
Accept? (y/n)
```

Note that, as our configuration space consists of variables of mixed units, the norm computed in (23) has been projected onto the Cartesian plane. A second prompt is given:

```
When within 1.44 meters of station_2, never set environment variable s2_occupied to True.
Accept? (y/n)
```

At this point, should the user accept both revisions, a new counterstrategy is synthesized containing no deadlock states. The user is prompted again:

```
Livelock revisions found.
Accept? (y/n)
```

This time, the specification is realizable if the user accepts and the resulting execution for the controller is as shown in Figure 8. The trajectories pictured in the figure represent evolutions of the continuous nonlinear system when commanded by the synthesized controller. Forward integration is applied to solve the equations of motion using an integration step size of 0.001 sec. Here $\varepsilon = 0.3$, and the regions are inflated to the extent indicated by the gray border. Note that the system in the figure infinitely often visits the three regions and is able to react to a change in the environment. In Figure 8a, the system avoids the region *station_1* when *s1_occupied* becomes True, since this does not happen within 1.32 m of *station_1*. A similar result is seen in Figure 8b. These behaviors are consistent with the intended behaviors encoded by the specification in Example 1.

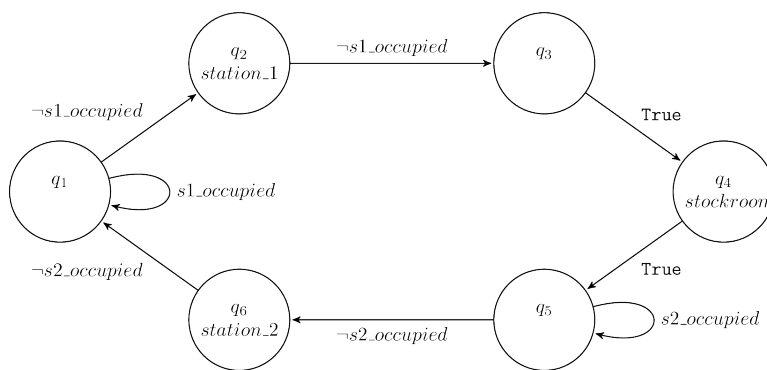


Fig. 7. Controller for φ in Example 1. Edges are labeled with the disjunction of assignments in \mathcal{X} that may be assumed for that transition.

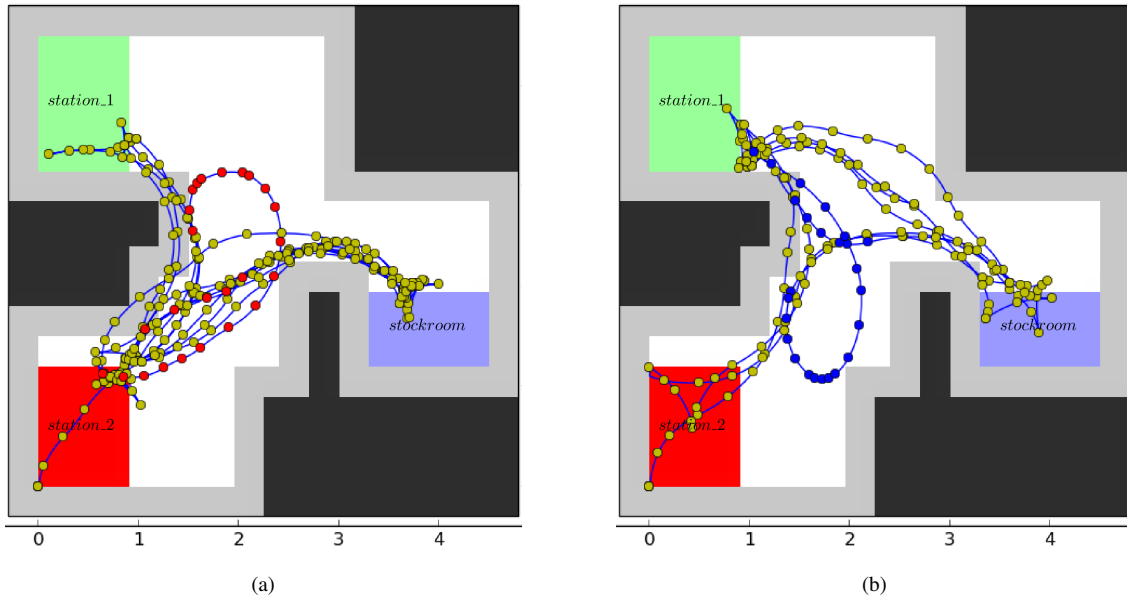


Fig. 8. Continuous trajectories for the nonlinear unicycle abstraction in a 5×5 workspace, initialized at the lower-left corner of the workspace. Regions *station_1*, *station_2* and *stockroom* are shown in green, red, and blue, respectively. Dots along the trajectory indicate position at each discrete time step (0.35 seconds). Color indicates the state of the environment (red: *s1_occupied*; blue: *s2_occupied*). (a) shows a trajectory when the *s1_occupied* sensor is activated. (b) shows a trajectory when the *s2_occupied* sensor is activated.

8. Conclusions

In this paper, we have described an automatic approach for synthesizing controllers for dynamical systems based on general specifications that are agnostic to the dynamics. We focus on the case where such specifications are transformed based on the discrete abstraction for a particular robot, and develop a framework for revising the specification in the case when the dynamics render the task unrealizable. We introduce a method for automatically generating a set of LTL formulas that, when added to the original specification, render it realizable. The approach features a mechanism for providing feedback to the user, giving him or her the freedom to accept or reject any such proposed formula. To facilitate interpretation of such formulas, the feedback given to the user is metric information that encapsulates the required restrictions on the environment and system behaviors due to the suggested revisions. Future work includes providing a means for suggesting a richer set of possible revisions to give as feedback to the user, thereby offering him or her a multiplicity of possible options to apply (e.g. trading off modifying the robot’s behavior vs. restricting the environment). Such an extension will involve mining more complex formulas from the synthesis game and automatically translating such formulas into easy-to-understand explanations.

Acknowledgement

The authors thank Salar Moarref, Ufuk Topcu, and Rajeev Alur for insightful discussions relating to synthesis of counterstrategy-based environment revisions.

References

- R. Alur, et al. (2013). ‘Counter-strategy guided refinement of GR(1) temporal logic specifications’. In *Formal Methods in Computer-Aided Design (FMCAD 2013)*, pp. 26–33.

- D. Angeli & E. D. Sontag (1999). ‘Forward completeness, unboundedness observability, and their Lyapunov characterizations’. *Systems & Control Letters* **38**(4):209–217.
- A. Bhatia, et al. (2010). ‘Sampling-based motion planning with temporal goals’. In *IEEE International Conference on Robotics and Automation (ICRA 2010)*, pp. 2689–2696. IEEE.
- R. Bloem, et al. (2012). ‘Synthesis of reactive (1) designs’. *Journal of Computer and System Sciences* **78**(3):911–938.
- J. DeCastro & H. Kress-Gazit (2014). ‘Synthesis of Nonlinear Continuous Controllers for Verifiably-Correct High-Level, Reactive Behaviors’. *International Journal of Robotics Research* Accepted.
- G. E. Fainekos (2011). ‘Revising Temporal Logic Specifications for Motion Planning’. In *Proceedings of the IEEE Conference on Robotics and Automation*.
- G. E. Fainekos, et al. (2009). ‘Temporal logic motion planning for dynamic robots’. *Automatica* **45**(2):343 – 352.
- G. E. Fainekos, et al. (2006). ‘Translating Temporal Logic to Controller Specifications’. In *Proc. of the 45th IEEE Conf. on Decision and Control (CDC 2006)*, pp. 899–904.
- G. E. Fainekos & G. J. Pappas (2009). ‘Robustness of temporal logic specifications for continuous-time signals’. *Theoretical Computer Science* **410**(42):4262–4291.
- A. Girard, et al. (2010). ‘Approximately bisimilar symbolic models for incrementally stable switched systems’. *Automatic Control, IEEE Transactions on* **55**(1):116–126.
- M. Kloetzer & C. Belta (2008). ‘Dealing with Nondeterminism in Symbolic Control’. In M. Egerstedt & B. Mishra (eds.), *Hybrid Systems: Computation and Control, 11th International Workshop (HSCC 2008)*, vol. 4981 of *Lecture Notes in Computer Science*, pp. 287–300. Springer.
- R. Könighofer, et al. (2009). ‘Debugging formal specifications using simple counterstrategies’. In *Proceedings of 9th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2009*, pp. 152–159.
- H. Kress-Gazit, et al. (2009). ‘Temporal Logic based Reactive Mission and Motion Planning’. *IEEE Transactions on Robotics* **25**(6):1370–1381.
- H. Kress-Gazit, et al. (2011). ‘Correct, Reactive Robot Control from Abstraction and Temporal Logic Specifications’. *Special Issue of the IEEE Robotics and Automation Magazine on Formal Methods for Robotics and Automation* **18**(3):65–74.
- W. Li, et al. (2011). ‘Mining assumptions for synthesis’. In *9th IEEE/ACM International Conference on Formal Methods and Models for Codesign, MEMOCODE 2011*, pp. 43–50.
- W. Li, et al. (2014). ‘Synthesis for Human-in-the-Loop Control Systems’. In *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014*, pp. 470–484.
- J. Liu & N. Ozay (2014). ‘Abstraction, Discretization, and Robustness in Temporal Logic Control of Dynamical Systems’. In *Proc. of the 17th Int. Conf. on Hybrid Systems: Computation and Control (HSCC’14) (to appear)*.
- J. Liu, et al. (2013). ‘Synthesis of Reactive Switching Protocols From Temporal Logic Specifications’. *IEEE Trans. Automat. Contr.* **58**(7):1771–1785.
- J. Liu, et al. (2012). ‘Reactive controllers for differentially flat systems with temporal logic constraints’. In *Proc. of the 51st IEEE Conf. on Decision and Control (CDC 2012)*, pp. 7664–7670.
- M. Maly, et al. (2013). ‘Iterative Temporal Motion Planning for Hybrid Systems in Partially Unknown Environments’. In *ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, pp. 353–362, Philadelphia, PA, USA. ACM, ACM.
- I. M. Mitchell, et al. (2005). ‘A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games.’. *IEEE Trans. Automat. Contr.* **50**(7):947–957.
- G. Pola, et al. (2008). ‘Approximately bisimilar symbolic models for nonlinear control systems’. *Automatica* **44**(10):2508–2516.
- V. Raman & H. Kress-Gazit (2013). ‘Explaining Impossible High-Level Robot Behaviors’. *IEEE Transactions on Robotics* **29**(1):94–104.
- V. Raman & H. Kress-Gazit (2013). ‘Towards Minimal Explanations of Unsynthesizability for High-Level Robot Behaviors’. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2013)*.
- G. Reißig (2011). ‘Computing abstractions of nonlinear systems’. *IEEE Transactions on Automatic Control* **56**(11):2583–2598.

- P. Tabuada & G. J. Pappas (2006). 'Linear Time Logic Control of Discrete-Time Linear Systems'. *IEEE Trans. Automat. Contr.* **51**(12):1862–1877.
- C. J. Tomlin, et al. (2000). 'A Game Theoretic Approach to Controller Design for Hybrid Systems'. *Proceedings of IEEE* **88**:949–969.
- J. Tumova, et al. (2010). 'A symbolic approach to controlling piecewise affine systems'. In *49th IEEE Conference on Decision and Control (CDC)*, pp. 4230–4235.
- M. Y. Vardi (1996). 'An automata-theoretic approach to linear temporal logic'. In *Logics for concurrency*, pp. 238–266. Springer.
- E. M. Wolff, et al. (2013). 'Automaton-Guided Controller Synthesis for Nonlinear Systems with Temporal Logic'. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS 2013)*.
- T. Wongpiromsarn, et al. (2010). 'Receding Horizon Control for Temporal Logic Specifications'. In *Proc. of the 13th Int. Conf. on Hybrid Systems: Computation and Control (HSCC'10)*.
- B. Yordanov, et al. (2012). 'Temporal Logic Control of Discrete-Time Piecewise Affine Systems'. *Automatic Control, IEEE Transactions on* **57**(6):1491–1504.
- M. Zamani, et al. (2012). 'Symbolic models for nonlinear control systems without stability assumptions'. *IEEE Transactions on Automatic Control* **57**(7):1804–1809.