

Enhancing tolerance to unexpected jumps in GR(1) games

Sumanth Dathathri

Scott C. Livingston

Richard M. Murray*

ABSTRACT

When used as part of a hybrid controller, finite-memory strategies synthesized from linear-time temporal logic (LTL) specifications rely on an accurate dynamics model in order to ensure correctness of trajectories. In the presence of uncertainty about the underlying model, there may exist unexpected trajectories that manifest as unexpected transitions under control of the strategy. While some disturbances can be captured by augmenting the dynamics model, such approaches may be conservative in that bisimulations may fail to exist for which strategies can be synthesized. In this paper, we consider games of the GR(1) fragment of LTL, and we characterize the tolerance of hybrid controllers to perturbations that appear as unexpected jumps (transitions) to states in the discrete strategy part of the controller. As a first step, we show robustness to certain unexpected transitions that occur in a finite manner, i.e., despite a certain number of unexpected jumps, the sequence of states obtained will still meet a stricter specification and hence the original specification. Additionally, we propose algorithms to improve robustness by increasing tolerance to additional disturbances. A robot gridworld example is presented to demonstrate the application of the developed ideas and also to perform empirical analysis.

1. INTRODUCTION

The ability of strategies synthesized from formal specifications to be tolerant to unexpected events such as

*S. Dathathri and R. M. Murray are with the California Institute of Technology, Pasadena, CA, USA. S. C. Livingston is independent. Email addresses are sdathath@caltech.edu and {slivingston,murray}@cds.caltech.edu.

disturbances or failures is important, especially for safety-critical applications. Reactive strategies synthesized to meet temporal logic objectives are not error resilient by default. Even with non-critical disturbances that were not accurately modeled during synthesis, no guarantees can be provided about satisfaction of the temporal-formula used for synthesis. In some cases it may not be possible to synthesize a strategy if these perturbations are fully modelled through the dynamics.

In this paper, we make progress towards enhancing the tolerance of strategies synthesized to satisfy specifications in the generalized reactivity(1) (GR(1)) fragment of linear-temporal logic (LTL) [11, 12]. GR(1) formulae are considered because they are quite expressive in terms of temporal properties captured, yet symbolic synthesis algorithms are possible with relatively low computational complexity [3, 8, 10]. The first result we show is that by refining a strategy synthesized to satisfy a GR(1) formula, a strategy that is winning against a stricter formula can be generated and is robust to certain unexpected events. Then, exploiting this tolerance, we propose multiple algorithms that combine separately synthesized strategies to form a single robust winning strategy. It is often desired that the system can recover from these unexpected events and this be done without resynthesizing the entire strategy again. In this regard, we propose an approach which lets us recover from these unexpected events without a complete resynthesis.

Understanding the response of systems to perturbations has been extensively studied in control theory and more recently, for reactive controllers and their synthesis. In [9, 13], the robustness considered is in terms of bounded input-output deviation. This relates directly to the prevalent notion in control for robustness [16], where controllers are designed to ensure bounded disturbances lead to bounded deviations from nominal-behavior for the system. In the current work, the tolerance to disturbances is in the form of satisfaction of a formula representing the desired system behavior. In [2], the effect of disturbances on system behavior is quanti-

fied and the focus on synthesizing robust systems that degrade gracefully - smallest number of system failures possible but not primarily directed at GR(1) specifications. Some existing work on notions of robustness in terms of satisfaction or violation of a formula can be found in [1, 5, 14].

In [6], they use a similar underlying idea to the one in the current work to re-synthesize a strategy against a new GR(1) formula that is more robust. In scenarios where unforeseen perturbations occur when the controller is implemented on the hybrid system, the results presented in our paper allow for continued execution with guarantees in terms of formula satisfaction. In [15], motivated by a similar idea, they propose an approach which uses the saved results from an intermediate step during synthesis to construct sequences of control actions that can tolerate violations of environmental safety assumptions. However, the work presented here differs in that the perturbations tolerated are more general. Additional guarantees are provided for instances where there is a failure on the system end. This has other potential implications as discussed in Section 6. Additionally, an approach to recover from both system failures and environmental assumption violations when possible is proposed that does not involve enumeration of an entire GR(1) strategy as in [15].

In summary, the main contributions of this work are the following: 1) to characterize the inherent tolerance of GR(1) strategies to unexpected events; 2) to propose approaches to refine GR(1) strategies to augment their tolerance to unexpected perturbations and prove that they satisfy a stricter temporal formula; 3) to quantify empirically the cost of augmenting the tolerance using the proposed schemes.

2. PRELIMINARIES

Let Σ be a finite set. The power set of Σ (i.e., the set of all subsets of Σ) is denoted by $\mathcal{P}(\Sigma)$. The set of all finite strings formed from concatenating elements of Σ is denoted by Σ^* , which is known as the Kleene closure [7]. The set of all countably infinite strings of Σ is Σ^ω . In this paper, a subscript notation is used, e.g., $\sigma_0\sigma_1\sigma_2\cdots\sigma_n \in \Sigma^*$, but observe that infinite strings can also be regarded as functions of the natural numbers \mathbb{N} into Σ .

Let AP_{env} be a set of input atomic propositions, and AP_{sys} be a set of output atomic propositions such that $\text{AP}_{\text{env}} \cap \text{AP}_{\text{sys}} = \emptyset$. A state s is an assignment of **True** and **False** to the atomic propositions in $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}}$. We use subset notation to indicate states and thus, for brevity, introduce $\Sigma = \mathcal{P}(\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}})$.

A *nondeterministic finite-memory strategy* is a pair (f, m_0) together with a finite set M , where $m_0 \in M$

and

$$f : M \times \mathcal{P}(\text{AP}_{\text{env}}) \rightarrow \mathcal{P}(M \times \mathcal{P}(\text{AP}_{\text{sys}})) \quad (1)$$

is a function. Intuitively the set M represents the memory of the strategy. At each move, a new output is given depending on the input and the current memory value. As part of the move, a memory value is selected. Since we are only concerned with finite-memory strategies in this paper, we refer to them as *strategies*. The set of input-output sequences that may occur under f is defined as

$$\text{Plays}(f) = \{\sigma \in \Sigma^\omega \mid \exists m \in M^\omega. \forall k \geq 0. (m_{k+1}, \sigma_k \cap \text{AP}_{\text{sys}}) \in f(m_k, \sigma_k \cap \text{AP}_{\text{env}})\}, \quad (2)$$

where every $m \in M^\omega$ has the same first element, m_0 . Elements of $\text{Plays}(f)$ are referred to as *plays*. The set of prefixes that may be extended into a play is

$$\text{Pref}(f) = \{\sigma \in \Sigma^* \mid \exists \alpha \in \Sigma^\omega. \sigma\alpha \in \text{Plays}(f)\}. \quad (3)$$

We describe specifications for these strategies in linear-time temporal logic (LTL) [12]. LTL formulae over propositions $(\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}})$ are evaluated over positions i in $\sigma = \sigma_0\sigma_1\cdots \in \Sigma^\omega$. In addition to the Boolean operators, the standard LTL operators \Box (always), \Diamond (eventually) and \bigcirc (next) are used here for the specification.

A finite-memory strategy (f, m_0) is said to be

- *input-enabled* if and only if for every $\sigma^{\text{env}} \in \mathcal{P}(\text{AP}_{\text{env}})^\omega$, there exists $\sigma \in \text{Plays}(f)$ such that $\sigma_k^{\text{in}} = \sigma_k \cap \text{AP}_{\text{env}}$ for $k \geq 0$.
- *deterministic* if and only if $|f(m, s)| \leq 1$ for all $(m, s) \in M \times \mathcal{P}(\text{AP}_{\text{env}})$.
- a *realization of* an LTL formula φ if and only if (f, m_0) is input-enabled and $\text{Plays}(f) \subseteq L(\varphi)$, where $L(\varphi)$ is the language of φ . This is to say, for every $\sigma \in \text{Plays}(f)$, $\sigma \models \varphi$.

Deterministic strategies imply a function from state sequences to memory sequences.

REMARK 1. *Let (f, m_0) be a deterministic strategy. For each $\sigma \in \text{Plays}(f)$, there exists a unique $m \in M^\omega$ satisfying $f(m_k, \sigma_k \cap \text{AP}_{\text{env}}) = (m_{k+1}, \sigma_k \cap \text{AP}_{\text{sys}})$ for $k \geq 0$.*

It follows from the remark that a sequence of inputs determines precisely one output sequence for a deterministic strategy.

A GR(1) formula is an LTL formula of the form

$$\Theta^{\text{env}} \wedge \Box \rho^{\text{env}} \wedge \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right) \implies \Theta^{\text{sys}} \wedge \Box \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \Box \Diamond \psi_k^{\text{sys}} \right), \quad (4)$$

where Θ^{env} is a state formula (i.e., without temporal operators) that is a function of AP_{env} , Θ^{sys} is a state formula that is a function of AP_{sys} , and all $\psi_j^{\text{env}}, \psi_k^{\text{sys}}$ subformulae are functions of $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}}$ and also without temporal operators. The subformula ρ^{env} is a function of $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}} \cup \bigcirc \text{AP}_{\text{env}}$, where

$$\bigcirc \text{AP}_{\text{env}} = \{\bigcirc x \mid x \in \text{AP}_{\text{env}}\}.$$

Except for \bigcirc operators appearing as subformulae from $\bigcirc \text{AP}_{\text{env}}$, there are no other temporal operators in ρ^{env} . Finally, ρ^{sys} is defined similarly to ρ^{env} but as a function of $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}} \cup \bigcirc \text{AP}_{\text{env}} \cup \bigcirc \text{AP}_{\text{sys}}$.

To facilitate working with (20), and in particular the subformulae ρ^{env} and ρ^{sys} , we extend the semantics of the operator \models for finite strings. For a finite string γ , by γ_{-1} we refer to the last element of γ . In other words, $\gamma_{-1} = \gamma_{|\gamma|-1}$. Let $\sigma \in \Sigma^*$. To further simplify notation, the superscripts sys and env for $s \in \Sigma$ indicate projections on to AP_{sys} and AP_{env} respectively. That is, given s , $s^{\text{sys}} := s \cap \text{AP}_{\text{sys}}$ and $s^{\text{env}} := s \cap \text{AP}_{\text{env}}$. Define

$$\sigma \models \rho \iff \sigma\alpha \models \rho \text{ for any } \alpha \in \Sigma^\omega, \quad (5)$$

where ρ is any Boolean formula that is a function of $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}} \cup \bigcirc \text{AP}_{\text{env}} \cup \bigcirc \text{AP}_{\text{sys}}$. Because at most one \bigcirc operator binds to each atomic proposition, it follows that only σ_0, σ_1 determine whether the formula is satisfied.

Let φ be a GR(1) formula, i.e., be of the form (20), and let (f, m_0) be a finite-memory strategy.

DEFINITION 2. *A state $s \in \Sigma$ is said to be φ -reachable under (f, m_0) if and only if there exists $\sigma \in \text{Plays}(f)$ such that for some $k \geq 0$,*

$$\sigma_k = s, \quad (6)$$

$$\sigma \models \Theta^{\text{env}}, \quad (7)$$

$$\sigma_{j:(j+1)} \models \rho^{\text{env}} \text{ for } j < k - 1. \quad (8)$$

The set of all states that are φ -reachable under (f, m_0) is denoted by $I_\varphi(f_0, m_0)$.

The finite-memory strategy (f, m_0) is said to be a *strict realization* of (or to *strictly realize*) the GR(1) formula φ if it is a realization of φ and for all $\sigma \in \text{Plays}(f)$

$$\sigma \models \Theta^{\text{env}} \implies \sigma \models \Theta^{\text{sys}} \quad (9)$$

$$\sigma \models (\exists \rho^{\text{env}} \implies \exists \rho^{\text{sys}}). \quad (10)$$

Intuitively, strict realizability ensures that blocking of an environment liveness condition when the other assumptions are met only occurs when the system is following transition rules. Here, \exists is the ‘historically’ LTL operator whose semantics are as defined in [10].

3. PROBLEM FORMULATION

3.1 LTL on state-memory pairs

A basic idea behind the methods presented in later sections is to begin with a given finite-memory strategy and construct another one that is more robust. The construction involves re-using ‘pieces’ of the given strategy, such as memory-values and transitions. As such, the memory values occurring during each play become important for reasoning about correctness and resilience to perturbations. Thus motivated, the semantics of \models for LTL as invoked in Section 2 is extended to handle memory values of strategies. Let (f, m_0) be a strategy that strictly realizes φ , and has set M of memory values. Following the notation used in the previous section, let $\Sigma = 2^{\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}}}$ be the set of game states. Define a set of state-memory pairs,

$$\bar{\Sigma} = \Sigma \times M.$$

An element $\bar{\sigma} \in \bar{\Sigma}^\omega$ is thus a sequence of state and memory pairs, i.e., there is $\sigma \in \Sigma^\omega$ and $m \in M^\omega$ such that for $k \geq 0$, $\bar{\sigma}_k = (\sigma_k, m_k)$. Given an LTL formula ϕ that is in terms of $\text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}}$ and a variable that takes values in M , the operator \models can thus be interpreted on $\bar{\sigma}$.

3.2 Memory sequences associated with plays

Because properties will be expressed in terms of sequences of states and memory values that occur during plays, a means for obtaining memory values that can occur in the strategy during a given play is needed. For this purpose, define the function Mem^f on $\text{Pref}(f) \cup \text{Plays}(f)$ as follows. Let $\sigma \in \text{Pref}(f) \cup \text{Plays}(f)$. Define $\sigma_k^{\text{sys}} := \sigma_k \cap \text{AP}_{\text{sys}}$ and $\sigma_k^{\text{env}} := \sigma_k \cap \text{AP}_{\text{env}}$. If $\sigma \in \text{Pref}(f)$, then define

$$\text{Mem}^f(\sigma) = \{m \in M^* \mid \forall k \geq 0, k < |\sigma|. \\ (m_{k+1}, \sigma_k^{\text{sys}}) \in f(m_k, \sigma_k^{\text{env}})\}. \quad (11)$$

If $\sigma \in \text{Plays}(f)$, then define

$$\text{Mem}^f(\sigma) = \{m \in M^\omega \mid \forall k \geq 0. \\ (m_{k+1}, \sigma_k^{\text{sys}}) \in f(m_k, \sigma_k^{\text{env}})\}. \quad (12)$$

Observe that $\text{Pref}(f) \cap \text{Plays}(f) = \emptyset$ (recall (2) and (3)), hence Mem^f is well-defined.

3.3 Problem of recovery from perturbations

Though there are many distinct notions of robustness for control systems, a common theme is tolerance to deviation from nominal plant behavior. Note that the ‘plant’ includes actuators and sensors on the robot itself. Thus in the context of reactive synthesis, deviance can also arise on the side of the controlled system, not only the adversarial environment.

We begin by introducing a few definitions. Define the set $I_\varphi^M(f_0, m_0) \subseteq \bar{\Sigma}$ as below:

$$I_\varphi^M(f, m_0) := \{(s, \bar{m}) \mid s \in I_\varphi(f, m_0), \bar{m} = m_{|\sigma s|}, \\ \sigma \in \Sigma^*, m \in M^*, \\ \sigma s \in \text{Pref}(f) \wedge \\ m \in \text{Mem}^f(\sigma s) \wedge \\ \forall k < |\sigma|. \sigma_k \in I_\varphi(f, m_0)\}$$

$I_\varphi^M(f, m_0)$ contains the set of φ -reachable states along with the corresponding memory locations. Given that we have the set of φ -reachable states, we need to extract all paths to these φ -reachable states that have only φ -reachable states to get all possible combinations of memory, φ -reachable states to which the system can be perturbed to. Given a φ -reachable state s , we find a prefix (σs) that ends with this state and this prefix has only φ -reachable states (Condition: $\forall k < |\sigma|. \sigma_k \in I_\varphi(f, m_0)$). Once this prefix has been found, we generate the memory sequence in accordance with (f, m_0) to get the corresponding memory at this state in the strategy [Condition: $\forall k < |\sigma|. ((\sigma s)_{k+1} \cap \text{AP}_{\text{sys}}, m_{k+1}) = f((\sigma s)_{k+1} \cap \text{AP}_{\text{env}}, m_k)$]

Using the definitions above, a function is constructed that indicates feasible transitions (in a sense to be made precise) that are not in the given strategy f . This function is crucial for studying perturbations and applications of a given strategy that are extrinsic to its original semantics. To this end, define the set of all state-memory pairs that can be reached by some play of f , i.e.,

$$I(f) = \{(s, w) \mid \exists \sigma \in \text{Plays}(f) \exists m \in \text{Mem}^f(\sigma) : \\ \exists k : s = \sigma_k \wedge w = m_{k+1}\}. \quad (13)$$

Using this definition of reachable state-memory pairs, we introduce the notion of a ‘perturbation’ for a strategy (f, m_0) that realizes φ .

DEFINITION 3. *A perturbation for a controller implementing the finite-memory strategy (f, m_0) occurs when the system transitions from a state-memory pair $(s, w) \in I(f)$ to a state s' in Σ such that $(w', s' \cap \text{AP}_{\text{sys}}) \notin f(w, s' \cap \text{AP}_{\text{env}}). \forall w' \in M \vee ss' \not\models \rho^{\text{env}}$.*

The condition $ss' \models \neg \rho^{\text{env}}$ corresponds to the environment violating an assumption on its behavior. And $(w', s' \cap \text{AP}_{\text{sys}}) \notin f(w, s' \cap \text{AP}_{\text{env}}). \forall w' \in M$ corresponds to a disturbance during the application of an output action. In this work we propose approaches that enable recovery from certain such perturbations.

Define the function ExtTs on $I(f)$ as follows. Let

$$(s, w) \in I(f).$$

$$\text{ExtTs}(s, w) = \{(s', w') \mid s' \in \Sigma, w' \in M.$$

$$(w', s') \in I_\varphi^M(f, m_0) \wedge \\ ss' \models \rho^{\text{sys}} \wedge \\ \left((w', (s')^{\text{sys}}) \notin f(w, (s')^{\text{env}}) \vee ss' \models \neg \rho^{\text{env}} \right)\}.$$

The function ExtTs maps a state-memory pair (s, m) in a strategy to all state-memory pairs (s', m') that occur in $I_\varphi^M(f, m_0)$ that are not immediate successors according to (f, m_0) or violate environmental safety [Condition: $(w', s' \cap \text{AP}_{\text{sys}}) \notin f(w, s' \cap \text{AP}_{\text{env}}) \vee ss' \models \neg \rho^{\text{env}}$]. Additional constraints are imposed, a transition between the state-memory pair (s, m) and a state-memory pair $[(s', m') \in \text{ExtTs}(s, m)]$ to which it is mapped to does not violate ρ^{sys} i.e $ss' \models \rho^{\text{sys}}$.

3.4 Refinement of a strategy

Given φ and (f, m_0) we construct a non-deterministic strategy (\bar{f}, m_0) with $\bar{f} : \mathcal{P}(\text{AP}_{\text{env}})^\omega \times M \rightarrow \mathcal{P}(\bar{\Sigma}^\omega)$. Because the strategy is non-deterministic, for a given sequence of inputs, there is a set of state-memory sequences generated.

$$\text{Plays}(\bar{f}) = \left\{ \sigma \in \Sigma^\omega \mid \exists m \in M^\omega. (\forall k > 0. \\ (m_{k+1}, \sigma_k^{\text{sys}}) \in f(m_k, \sigma_k^{\text{env}}) \\ \vee (\sigma_k, m_{k+1}) \in \text{ExtTs}(\sigma_{k-1}, m_k) \right\} \\ \wedge (m_1, \sigma_0^{\text{sys}}) \in f(m_0, \sigma_0^{\text{env}}) \quad (14)$$

Consider $\sigma \in \text{Plays}(\bar{f})$ and $m \in M^\omega$ such that it satisfies the conditions in (14). Observe that the successor memory-state pair (m', s') in $\text{Plays}(\bar{f})$ depends on the current memory-state pair (m, s) and the next input $s_{\text{input}} = s' \cap \text{AP}_{\text{env}}$. By augmenting the memory-values with elements from Σ , new memory-values can be created that \bar{f} can use as memory inputs. Given a memory-state pair and an input in $\mathcal{P}(\text{AP}_{\text{env}})$, \bar{f} can map to an output in $\mathcal{P}(\text{AP}_{\text{env}})$ and a memory-value in the augmented memory domain. This way $\text{Plays}(\bar{f})$ can be used to construct a strategy \bar{f} (in the conventional sense) that maps from a memory value and an input to a memory value and an output.

Let $\bar{\sigma} \in \bar{\Sigma}^\omega$ such that $\forall k \geq 0. \bar{\sigma}_k = (\sigma_k, m_{k+1})$. Notice that here m is the memory sequence corresponding to σ in accordance with \bar{f} .

To simplify notation, for $(s, m) \in \bar{\Sigma}$ introduce the indicator formula

$$\mathbf{1}_{(s, m)} := \bigwedge_{p \in s \cup \{m\}} p \quad \wedge \quad \bigwedge_{q \in \text{AP}_{\text{env}} \cup \text{AP}_{\text{sys}} \cup M / s \cup \{m\}} \neg q$$

. Define 1_{jump} as

$$1_{jump} := \bigwedge_{(s,m) \in \text{dom}(\text{ExtTs})} \left(\mathbf{1}_{(s,m)} \rightarrow \bigcirc \bigvee_{(s',m') \in \text{ExtTs}(s,m)} \mathbf{1}_{(s',m')} \right)$$

PROPOSITION 4. For all such $\bar{\sigma}$, the following holds:

$$\begin{aligned} \bar{\sigma} \models \Theta^{\text{env}} \wedge \square(\rho^{\text{env}} \vee 1_{jump}) \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right) \wedge \diamond \square \neg 1_{jump} \\ \implies \Theta^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right) \quad (15) \end{aligned}$$

This proposition arises as a special case of Proposition 5 where $n = 0$ in the setup before Proposition 5. Proposition 5 proposes an approach to combine multiple strategies to provide tolerance to perturbations, when there is just one strategy the current proposition results.

Intuitively, the practical significance of Proposition 4 is that, if there is a disturbance that causes an unexpected transition to some state that is φ -reachable in other plays and if there are only finitely many such disturbances, then execution of the finite-state machine can continue after an appropriate change of its internal state and still result in a correct input-output sequence. If ρ^{env} is violated during a particular transition between a state s and its successor s' , i.e. $ss' \not\models \rho^{\text{env}}$, and s' is φ -reachable, then a sequence of input-outputs that satisfies the system part of φ is still possible if there are finitely many such disturbances. It also allows for such disturbances during application of the output action to a hybrid system – where $(s' \cap \text{AP}_{\text{sys}}, m') \notin f(s' \cap \text{AP}_{\text{env}}, m)$. Here m is the current memory value and $s' \cap \text{AP}_{\text{env}}$ the input and $s' \cap \text{AP}_{\text{sys}}$ the ground truth system state after the application of the output action.

However, this result not does imply that all disturbances in which the system fails to transition to the desired state or the transition rule ρ^{env} is violated can be handled. Only perturbations to φ -reachable states that do not violate ρ^{sys} are tolerated.

The added $\diamond \square \neg 1_{jump}$ segment on the environment-assumption side in 15 ensures that the perturbations do not occur infinitely often.

4. AUGMENTING ROBUSTNESS

To illustrate the methods proposed in this section to augment robustness, we begin with a small deterministic example. Consider planar robotic motion planning on a 4×4 grid shown in Figure 1. The robot begins in the cell marked ‘I’ and has to visit the cells marked ‘G’ infinitely

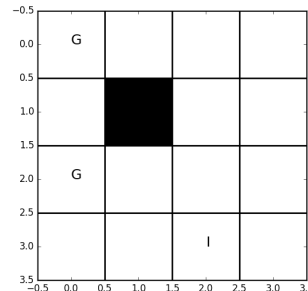


Figure 1: Gridworld setup

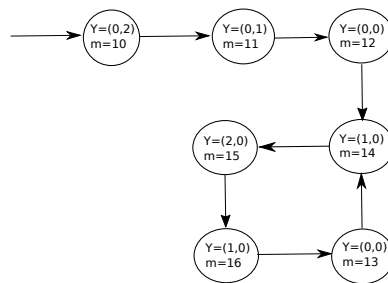


Figure 2: Strategy 1

often. The shaded cell indicates a wall the robot must not collide with, and the robot can transition to any of its non-diagonally adjacent cells.

For the system, two strategies that satisfy the safety and progress requirements were synthesized, but they differ in their initial poses on the gridworld. While the robot was executing Strategy 1 (in Figure 2), a perturbation causes the robot to transition from the cell $Y = (Y_r, Y_c) = (0, 1)$ with memory $m = 11$ to $(0, 2)$ where Y_r refers to the row position of the robot and Y_c refers to the column position. This transition is not desired according to Strategy 1 but is still safe for the system, i.e. ρ^{sys} is still satisfied by the move of the robot from $(0, 1)$ to $(0, 2)$. However, the strategy fails to predict what the sequence of actions from there should be such that φ can be satisfied. Proposition 5 guarantees that continued execution along strategy 2 from the state $(0, 2)$ with memory $m = 3$ would satisfy φ for the robot. Here, Strategy 2 (in Figure 3a) was synthesized to satisfy a specification similar to the one used to synthesize Strategy 1.

Now consider a scenario where it is not desirable to resynthesize or enumerate a new strategy. We wish to find a sequence of actions to a φ -reachable state visited by Strategy 1 such that continued execution along Strategy 1 satisfies φ . Proposition 6 guarantees the correctness of the algorithm to find such a sequence of actions. For this example, Strategy 3 (in Figure 3b) depicts one such path that was synthesized starting from

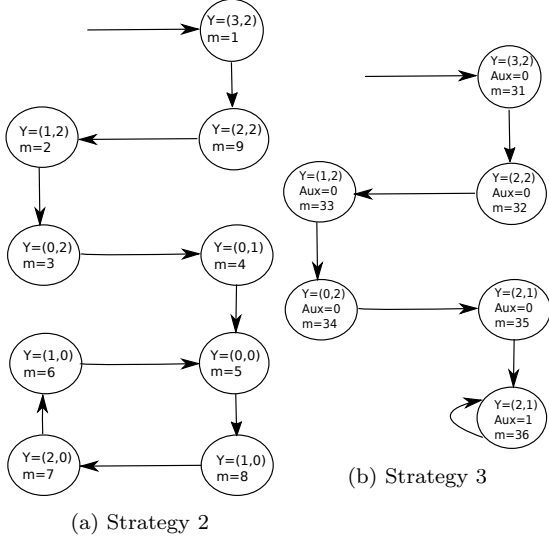


Figure 3: Synthesized strategies to augment robustness

(3, 2) reconnecting to a φ -reachable state. The reader must note that the absence of an adversary in this example makes the recovery process trivial, but solving a reachability game for general GR(1) formulae so as to reconnect to a φ -reachable state in the original strategy is more involved.

4.1 Combining Multiple Strategies

In this section, the intuition from Section 3.4 is used and a more general proposition is presented and proved. The results in this section allow for the concatenation of multiple strategies synthesized with formulae differing in the initial condition.

Let φ_0 be a GR(1) formula, and let $(g_0, m_0) : M_0 \times \mathcal{P}(\text{AP}_{\text{env}}) \rightarrow \mathcal{P}(M_0 \times \mathcal{P}(\text{AP}_{\text{sys}}))$ be a strategy that realizes φ_0 . Let $\eta_1, \eta_2, \dots, \eta_n$ be the additional states in Σ that the strategy must visit to provide additional robustness. For $i \in \{1, 2, \dots, n\}$ where $n < \infty$, define Θ_i^{env} as a Boolean formula which is **True** for a state s in Σ if and only if $s \cap \text{AP}_{\text{env}} = \eta_i^{\text{env}}$. Similarly, define Θ_i^{sys} as a Boolean formula which is **True** for a state s in Σ if and only if $s \cap \text{AP}_{\text{sys}} = \eta_i^{\text{sys}}$.

Then, construct a set of strategies (g_i, m_0^i) such that $\forall i \in \{1, 2, \dots, n\}$, (g_i, m_0^i) strictly realizes φ_i , where

$$\varphi_i = \Theta_i^{\text{env}} \wedge \square \rho^{\text{env}} \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right) \implies \Theta_i^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right). \quad (16)$$

To simplify notation, without loss of generality let $g_i : M_i \times \mathcal{P}(\text{AP}_{\text{env}}) \rightarrow \mathcal{P}(M_i \times \mathcal{P}(\text{AP}_{\text{sys}}))$ such that $\forall i, j \in \{0, 1, \dots, n\}, i \neq j \rightarrow M_i \cap M_j = \emptyset$.

Let $\bar{M} = \bigcup_{i=0,1,2,\dots,n} M_i$. Define a partial function

Algorithm 1 Implements controller based on Section 4.1

Input: • Finite-memory strategies $(g_i, m_0^i) \forall i \in \{0, 1, 2, \dots, n\}$,

• Sequence of inputs $\sigma^{\text{env}} \in \text{AP}_{\text{env}}^\omega$,

• A system to which the sequence $\sigma^{\text{sys}} \in (2^{\text{AP}_{\text{env}}})^\omega$ can be applied to and its state $s \in \Sigma$ after the control action has been applied

• Set $I_{\varphi_i}^M(g_i, m_0^i)$ and function ExtTs

Output: Sequence of control actions $\sigma^{\text{sys}} \in \mathcal{P}(\text{AP}_{\text{sys}}^\omega)$

```

1: memory=m0, i=1, safety=1, l=0
2: (memoryNew, σ0sys) = Strategyf : (memory, σ0env)
3: while (True) do
4:   if (σi-1env, σi-1sys)(σienv) ⊨ ρenv and safety=1 then
5:     Choose (memoryNew, σisys) ∈ fl : (memory, σienv)
6:     Run: SafetyCheck
7:   else if safety=0 ∧ (σi-1env+, σi-1sys)(σienv) ⊨ +ρenv then
8:     if ∃j ≤ n, m ∈  $\bar{M}.(m, \sigma_{i-1}^{\text{env}} \cup \sigma_{i-1}^{\text{sys}}) \in I_{\varphi_j}^M(f_j, m_0^j)$  then
9:       Choose (memoryNew, σisys) ∈ fj(m, σienv)
10:      l = j
11:      Run: SafetyCheck
12:     else
13:       EXIT
14:     end if
15:   else if (σi-1env, σi-1sys)(σienv) ⊭ ρenv then
16:     if ∃j ≤ n, m ∈ Mj, s ∈  $\mathcal{P}(\text{AP}_{\text{sys}}). (m, \sigma_{i-1}^{\text{env}} \cup s) \in \text{ExtTs}(\text{memory}, \sigma_{i-1}^{\text{env}} \cup \sigma_{i-1}^{\text{sys}})$  then
17:       (memoryNew, σisys) = (m, s)
18:       l = j
19:       Run: SafetyCheck
20:     else
21:       EXIT
22:     end if
23:   end if
24:   i+=1, memory=memoryNew
25: end while

```

$\text{ExtTs} : \bar{\Sigma} \times \bar{M} \rightarrow 2^{i=0,1,2,\dots,n} I_{\varphi}^M(g_i, m_0^i)$ such that for $(s, w) \in I(g_i)$ for some $i \leq n$,

$\text{ExtTs}(s, w) = \{(s', w') \mid s' \in \Sigma, w' \in M_i, i \in \{0, 1, \dots, n\}$.

$(s', w') \in I_{\varphi_i}^M(g_i, m_0^i) \wedge$

$ss' \models \rho^{\text{sys}} \wedge$

$\left((w', s' \cap \text{AP}_{\text{sys}}) \notin g_i(w, s' \cap \text{AP}_{\text{env}}) \vee ss' \models \neg \rho^{\text{env}} \right)$

Algorithm 2 SafetyCheck

```

1: apply  $\sigma_i^{\text{sys}}$ , measure  $s$ 
2: if  $(\sigma_{i-1}^{\text{env}}, \sigma_{i-1}^{\text{sys}})s \not\models \rho^{\text{sys}}$  then
3:   EXIT
4: end if
5: if  $(\sigma_i^{\text{env}}, \sigma_i^{\text{sys}}) = s$  then
6:   safety=1
7: else
8:   safety=0,  $\sigma_i^{\text{env}} = s \cap \text{AP}_{\text{env}}$ ,  $\sigma_i^{\text{sys}} = s \cap \text{AP}_{\text{sys}}$ 
9: end if
  
```

Define $\bar{1}_{jump}$ as

$$\bar{1}_{jump} := \bigwedge_{(s,m) \in \text{dom}(\text{ExtTs})} \left(\mathbf{1}_{(s,m)} \rightarrow \bigcirc \bigvee_{(s',m') \in \text{ExtTs}(s,m)} \mathbf{1}_{(s',m')} \right)$$

4.1.1 Robust Strategy from combining strategies

Construct a non-deterministic strategy (\bar{g}, m_0) .

$$\begin{aligned} \text{Plays}(\bar{g}) = \{ \sigma \in \Sigma^\omega \mid \exists m \in \bar{M}. \forall k > 0. \exists i. \\ (m_{k+1}, \sigma_k^{\text{sys}}) \in g_i(m_k, \sigma_k^{\text{env}}) \\ \vee \left((\sigma_k, m_{k+1}) \in \text{ExtTs}(\sigma_{k-1}, m_k) \right) \\ \wedge (m_1, \sigma_0^{\text{sys}}) \in g_0(m_0, \sigma_0^{\text{env}}) \} \end{aligned} \quad (17)$$

Consider $\sigma \in \text{Plays}(\bar{g})$ and $m \in (\bar{M})^\omega$ that satisfy the conditions in (17). Let $\bar{\sigma} \in \bar{\Sigma}^\omega$ such that $\forall k \geq 0. \bar{\sigma}_k = (\sigma_k, m_{k+1})$. Here, m is the memory-sequence corresponding to σ generated in accordance with \bar{g} .

PROPOSITION 5. *For all such $\bar{\sigma}$, the following holds:*

$$\begin{aligned} \bar{\sigma} \models \Theta_0^{\text{env}} \wedge \square(\rho^{\text{env}} \vee \bar{1}_{jump}) \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right) \wedge \square \neg \bar{1}_{jump} \\ \implies \Theta_0^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right), \end{aligned} \quad (18)$$

A proof is provided in the appendix.

Algorithm 1 gives a formal description of the implementation of a controller on a ‘plant’ that exploits the non-determinism in the strategy \bar{g} to tolerate certain perturbations. Algorithm 2 measures the state of the system after the application of the output action. If ρ^{sys} is violated because of a disturbance, execution is terminated. When there is a system-actuation failure $\left[(s' \cap \text{AP}_{\text{env}}, w') \notin g_j(s' \cap \text{AP}_{\text{sys}}, w) \right]$, it is checked if

the resulting state s' corresponds to a transition permissible in accordance with \bar{g} . If so, the memory is reset to an appropriate memory such that $(s', w') \in I_{\varphi_i}^M$ and execution is continued. At a state s , if the next input $(s')^{\text{env}} \in \mathcal{P}(\text{AP}_{\text{env}})$ is such that $s(s')^{\text{env}} \models \neg \rho^{\text{env}}$ then we choose $(s')^{\text{sys}}$ as indicated by ExtTs. Here, we choose to continue execution from a state indicated by ExtTs as continuing execution along the original strategy would give no guarantees in terms of system behavior. By choosing to deviate in accordance with Proposition 5, we are still able to provide guarantees above system behavior. To summarize, from a set of strategies $\{g_i : i \leq n\}$ we proposed an algorithm that can tolerate a certain set of perturbations finitely many times.

4.2 Building patches to handle perturbations

This section proposes an approach to add ‘patches’ to a strategy that enable recovery from perturbations by finding a safe sequence of states back to the original strategy. As opposed to the previous section where complete strategies were synthesized, here only a recovery patch is synthesized. Let $(h_0, m_0^0) : M_0 \times \mathcal{P}(\text{AP}_{\text{env}}) \rightarrow \mathcal{P}(\text{AP}_{\text{sys}})(M_0 \times \mathcal{P}(\text{AP}_{\text{sys}}))$ be a finite-memory strategy that strictly realizes a GR(1) formula φ . For $i \in \{1, 2, \dots, n\}$, let η_i be a state in Σ such that $\eta_i \notin I(f, m_0)$. Define $\Theta_i^{\text{env}}, \Theta_i^{\text{sys}}$ corresponding to η_i as in Section 4.1. Define \mathcal{T}_{reach} as a Boolean formula that evaluates to True at a state s in Σ if and only if $s \in I_\varphi(f)$. Define φ_{reach}^i as below:

$$\begin{aligned} \varphi_{reach}^i := \left(\Theta_i^{\text{env}} \wedge \square \rho^{\text{env}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{env}} \right) \right) \implies \\ (\Theta_i^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \diamond \mathcal{T}_{reach}) \end{aligned}$$

For $i \leq n$ let $(h^i, m_0^i) : M_i \times \mathcal{P}(\text{AP}_{\text{env}}) \rightarrow \mathcal{P}(\text{AP}_{\text{sys}})(M_i \times \mathcal{P}(\text{AP}_{\text{sys}}))$ be a finite-memory strategy that strictly realizes (defined as for the GR(1) specification) φ_{reach}^i . Define \bar{M} as before and let $M_i \cap M_j = \emptyset$ for any $i, j \leq n$ with $i \neq j$. Note that φ_{reach}^i can be converted to a GR(1) formula by introducing an auxiliary variable as in [4]. Define $I_{\varphi_{reach}^i}^{reach}$ as

$$\begin{aligned} I_{\varphi_{reach}^i}^{M-reach} := \{ (s, \bar{w}) \mid s \in I_{\varphi_{reach}^i}^i(h^i, m_0^i), \bar{w} = m_{|\sigma|s}. \\ \sigma \in \Sigma^*, m \in M_i^*. \\ \sigma s \in \text{Pref}(h^i) \wedge \\ m \in \text{Mem}^{h^i}(\sigma s) \wedge \\ \forall k < |\sigma|. (\sigma_k \in I_{\varphi_{reach}^i}^i(h^i, m_0^i) \wedge \sigma_k \models \neg \mathcal{T}_{reach}) \} \end{aligned}$$

This definition is similar to that of I_φ^M except for the additional constraint where we require that all states in the prefix are such that they are not from $I_\varphi(h^0, m_0^0)$.

Algorithm 3 Executing controller for a patched strategy

Input: • GR(1) formula φ

- Finite-memory strategy (h^0, m_0) realizing φ and a set of strategies (h^l, m_0^l) for $l \in \{1, 2, \dots, n\}$ realizing φ_{reach}^i
- Sequence of inputs $\sigma^{\text{env}} \in \mathcal{P}(\text{AP}_{\text{env}})^\omega$
- System whose state $s \in \Sigma$ is measured after a control action ($\in 2^{\text{AP}_{\text{sys}}}$) has been applied
- Sets $I_\varphi(h^0), I_0^M(h_0, m_0^0) := I_\varphi^M(h^0, m_0^0)$, $I_l^M(h^l, m_0^l) := I_{\varphi_{reach}^l}^{M-reach}(h^l, m_0^l)$ for $l \in \{1, 2, \dots, n\}$

Output: Sequence of control actions $\sigma^{\text{sys}} \in \mathcal{P}(\text{AP}_{\text{sys}})^\omega$

```

1: memory=m0
2: i=1, l=0, safety=1, reached=1
3: (memoryNew, σ0sys) = Strategy h0 : (memory, σ0env)
4: while (True) do
5:   if (σi-1env, σi-1sys)(σienv) ⊨ ρenv and safety=1 and
     reached=1 then
6:     (memoryNew, σisys) = Strategy h0 : (memory,
       σienv)
7:     Run: SafetyCheck
8:   else if (σi-1env, σi-1sys)(σienv) ⊨ ρenv and safety=1
     and reached=0 then
9:     (memoryNew, σisys) = Strategy hl : (memory,
       σienv)
10:    if σisys ∪ σienv ∈ Iφ(h0, m00) then
11:      Choose (memoryNew, σisys) such that (σisys ∪
        σienv, memoryNew) ∈ I0M
12:      reached=1
13:    end if
14:    Run: SafetyCheck
15:   else if safety=0 ∧ (σi-1env, σi-1sys)(σienv) ⊨ ρenv
     then
16:     if ∃j ≤ n, m ∈ M̄.(m, σi-1env ∪ σi-1sys) ∈
       IjM(hj, m0j) then
17:       Choose (memoryNew, σisys) ∈ fj(m, σienv)
18:       l = j, if j > 0 then reached=0.
19:       Run: SafetyCheck
20:     else
21:       EXIT
22:     end if
23:   else if (σi-1env, σi-1sys)(σienv) ⊭ ρenv then
24:     if ∃j ≤ n, m ∈ M̄j, s ∈ P(APsys). (m, σienv ∪
       s) ∈ ExtTsreach(memory, σi-1env ∪ σi-1sys) then
25:       (memoryNew, σisys) = (m, s)
26:       l = j, if j > 0 then reached=0.
27:       Run: SafetyCheck
28:     else
29:       EXIT
30:     end if
31:   end if
32:   i+=1, memory=memoryNew
33: end while

```

This is to allow perturbation to states in sequences before \mathcal{T}_{reach} was satisfied. This is because if a sequence of states in $\text{Plays}(h^i)$ (for $i > 0$) satisfies $\Theta_i^{\text{env}} \wedge \Box \rho^{\text{env}} \wedge (\bigwedge_{k=1}^K \Box \diamond \psi_k^{\text{env}})$ then it satisfies $\diamond \mathcal{T}_{reach}$. Since we are synthesizing a patch back to the original strategy, jumping to a state in $I^{M-reach}$ ensures that if the environment satisfies the assumptions on safety and liveness, we will reach a state where \mathcal{T}_{reach} holds.

Next, define ExtTs_{reach} for $(s, w) \in I(h^i)$ for some $i \leq n$ as

$$\begin{aligned}
\text{ExtTs}_{reach}(s, w) := & \{(s', w') \mid s' \in \Sigma, w' \in M_i, \\
& i \in \{0, 1, \dots, n\}. \\
(i = 0) \implies & (s', w') \in I_\varphi^M(h^0, m_0^0) \wedge \\
(i > 0) \implies & (s', w') \in I_{\varphi_{reach}^i}^{M-reach} \wedge \\
& ss' \models \rho^{\text{sys}} \wedge \\
& \left((w', (s')^{\text{sys}}) \notin h^i(w, (s')^{\text{env}}) \vee ss' \models \neg \rho^{\text{env}} \right) \}
\end{aligned}$$

This again is similar to the earlier definition of ExtTs except that we distinguish between the case where $i = 0$ and $i > 0$. Define \bar{I}_{jump} as:

$$\begin{aligned}
\bar{I}_{jump} := & \bigwedge_{(s, m) \in \text{dom}(\text{ExtTs}_{reach})} \left(\mathbf{1}_{(s, m)} \right. \\
& \left. \rightarrow \neg \bigcirc \bigvee_{(s', m') \in \text{ExtTs}_{reach}(s, m)} \mathbf{1}_{(s', m')} \right)
\end{aligned}$$

4.2.1 'Patched' Robust Strategy

Construct a non-deterministic strategy \bar{h} defined as below:

$$\begin{aligned}
\text{Plays}(\bar{h}) = & \{ \sigma \in \Sigma^\omega \mid \exists m \in \bar{M}. \forall k > 0. \exists i. \\
& (i = 0) \implies \left((m_{k+1}, \sigma_k^{\text{sys}}) \in h^i(m_k, \sigma_k^{\text{env}}) \right. \\
& \left. \vee (\sigma_k, m_{k+1}) \in \text{ExtTs}_{reach}(\sigma_{k-1}, m_k) \right) \\
\wedge & (i > 0) \implies \left[(\sigma_k, m_{k+1}) \in \text{ExtTs}_{reach}(\sigma_{k-1}, m_k) \right. \\
& \left. \vee \left(\forall w \in M_0. (w, \sigma_k) \notin I_\varphi^M(h^0, m_0^0) \wedge \right. \right. \\
& \left. \left. (m_{k+1}, \sigma_k^{\text{sys}}) \in h^i(m_k, \sigma_k^{\text{env}}) \right) \right] \\
& \wedge (m_1, \sigma_0^{\text{sys}}) \in h^0(m_0, \sigma_0^{\text{env}}) \} \tag{19}
\end{aligned}$$

The definition of \bar{h} is similar to \bar{f} defined earlier, except for $i > 0$. The condition $(i > 0) \rightarrow \left[(\sigma_k, m_{k+1}) \in \right.$

$\text{ExtTs}(\sigma_{k-1}, m_k) \vee \left(\forall w \in M_0.(w, \sigma_k) \notin I_\varphi^M(h^0, m_0) \wedge (m_{k+1}, \sigma_k^{\text{sys}}) \in h^i(m_k, \sigma_k^{\text{env}}) \right)$ ensures that unless perturbed according to $\text{ExtTs}_{\text{reach}}$, on reaching a state in state in $I_\varphi(h^0)$, execution is continued in accordance with the strategy h^0 . Consider $\sigma \in \text{Plays}(\bar{h})$ and $m \in (\bar{M})^\omega$ that satisfy the conditions in (19). Let $\bar{\sigma} \in \bar{\Sigma}^\omega$ such that $\forall k \geq 0. \bar{\sigma}_k = (\sigma_k, m_{k+1})$. Here, m is the memory-sequence corresponding to σ generated in accordance with \bar{h} .

PROPOSITION 6. *For all such $\bar{\sigma}$, the following holds:*

$$\bar{\sigma} \models \Theta_0^{\text{env}} \wedge \square(\rho^{\text{env}} \vee \bar{1}_{\text{jump}}) \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right) \wedge \square \neg \bar{1}_{\text{jump}} \\ \implies \Theta_0^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right), \quad (20)$$

A proof is provided in the appendix.

Algorithm 3 formally describes the implementation of a controller based on \bar{h} . The controller makes use of the non-determinism to recover from perturbations to the patches. Execution is begun along the strategy h^0 till a perturbation sets in. The controller attempts to recover using a patch (when perturbed onto a patch) if feasible and execution is continued along the patch to a state in $I_\varphi(h^0)$. Following this, execution along the original strategy is continued. Also, note the algorithm allows for jumps to φ -reachable states in the strategy itself (as does \bar{h}). The strategy also allows for perturbations from one patch to another during execution of the patch itself.

5. EXAMPLE IMPLEMENTATION AND ANALYSIS

Examples are implemented for the analysis of the techniques described in Section 4.1 and Section 4.2 for the task of planar robot motion planning in the environments similar to that shown in Figure 1. The robot is required to visit a set of locations infinitely often (progress-states). A moving obstacle whose behavior dynamics mimic those of the robot with different progress states and initial position is added to the setup.

5.1 Complexity for refinement

An empirical analysis of the computational costs involved in each of the approaches to augment robustness is presented here. The computations were performed on a 2.40GHz Quadcore machine with 16 GB of RAM. The experiment described below is repeated 50 times and the average synthesis times are presented in Table 1.

Random 5×5 gridworlds are generated with a wall density of 0.2. The moving obstacle and robot have two different progress-locations which they visit infinitely often and different initial positions. For each of the approaches perturbation points are chosen as below:

- **Multiple Strategy Approach:** A single perturbation point is chosen that is not visited by the initial strategy and a strategy is synthesized. The states visited by the new strategy are stored. A new perturbation point is chosen not visited by the earlier strategies. This repeated.
- **Patching:** The points are chosen as in the previous approach except only those states from the new strategy are stored that occurred before the trajectories reached the old set of φ -reachable states.

Approach	Coverage (unique states)	Mean time for synthesis(s)
New strategy from perturbed state	145.14	0.22
Patching	173.63	2.98

Table 1: Runtimes and unique states visited

The coverage i.e the number of unique states - robot, moving obstacle position combinations - visited by each strategy is also presented. It loosely characterizes the robustness for the concatenated-strategies as this count represents the φ -reachable states. Patching is implemented iteratively, with the visited states augmented in each patch. With iterative patching, the time for synthesis tends to decrease progressively with each patch for a given gridworld because the number of unique visited states tends to go up.

6. DISCUSSION AND CONCLUSION

This paper characterizes the inherent robustness of strategies synthesized to satisfy GR(1) formulae and also describes approaches to refine them to increase their tolerance to perturbations. It was shown that these refined strategies satisfy a stricter formula than the one used for synthesis. This tolerance is useful when the model is not exact for either the system behavior or the environment behavior. However, not all perturbations as defined in Section 3 can be tolerated through the approaches presented. Perturbations where ρ^{sys} is violated cannot be recovered from and perturbations to states that are not winning also cannot be recovered from using the approaches described here.

Another application of the results presented here is in the presence of noise in measurements. Reacting to environmental events in physical systems involves measurements about the environment, and in the presence

of noise, false inferences could be made under which the guarantees provided in terms of formula satisfaction would no longer hold. False inferences about an environmental event could alternatively be viewed as the system failing to apply the correct control action. Though ρ^{env} was satisfied, the false inference could result in the system applying an incorrect control action. Let $\gamma \in 2^{\text{AP}^{\text{env}}}$ be the inferred environmental event and the ground truth environmental event be $\mu \in 2^{\text{AP}^{\text{env}}}$. The system control action $s \in 2^{\text{AP}^{\text{sys}}}$ was decided based on $f(m_k, \gamma)$ where m_k is the current memory state. This can be viewed as a perturbation with $(m_{k+1}, s) \notin f(m_k, \mu)$ for any $m_{k+1} \in M$. If $s \cup \mu \in \Sigma$ is a state for which any of the above described approaches can be applied, the system can recover to satisfy the guarantees on system behavior.

Future Work. We plan to extend the framework built here to the case of infinitely many jumps. We also intend to develop a metric that would quantify the robustness added to a strategy through a given concatenation and prescribe approaches for refinement of strategies to make them more robust with optimal synthesis time/memory costs. Also, we plan to implement the approaches in Sections 4.1 and 4.2 using enumeration from a stored BDD computed during the original synthesis as opposed to resynthesis.

Acknowledgments

This work was partially supported by United Technologies Corporation and IBM, through the industrial cyber-physical systems (iCyPhy) consortium and by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA.

7. REFERENCES

- [1] R. Bloem, K. Chatterjee, K. Greimel, T. A. Henzinger, and B. Jobstmann. *Computer Aided Verification: 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, chapter Robustness in the Presence of Liveness, pages 410–424. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [2] R. Bloem, K. Greimel, T. A. Henzinger, and B. Jobstmann. Synthesizing robust systems. In *Formal Methods in Computer-Aided Design, 2009. FMCAD 2009*, pages 85–92, Nov 2009.
- [3] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa’ar. Synthesis of reactive(1) designs. *Journal of Computer and System Sciences*, 78:911–938, May 2012.
- [4] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Saar. Synthesis of reactive(1) designs. *J. Comput. Syst. Sci.*, 78(3):911–938, May 2012.
- [5] R. Ehlers. *NASA Formal Methods: Third International Symposium, NFM 2011, Pasadena, CA, USA, April 18-20, 2011. Proceedings*, chapter Generalized Rabin(1) Synthesis with Applications to Robust System Synthesis, pages 101–115. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [6] R. Ehlers and U. Topcu. Resilience to intermittent assumption violations in reactive synthesis. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control, HSCC ’14*, pages 203–212, New York, NY, USA, 2014. ACM.
- [7] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [8] Y. Kesten, N. Piterman, and A. Pnueli. Bridging the gap between fair simulation and trace inclusion. *Information and Computation*, 200:35–61, 2005.
- [9] R. Majumdar, E. Render, and P. Tabuada. Robust discrete synthesis against unspecified disturbances. In *Hybrid Systems: Computation and Control (HSCC)*, April 2011.
- [10] Z. Manna and A. Pnueli. A hierarchy of temporal properties. In *(PODC ’90) Proceedings of the ninth annual ACM Symposium on Principles of Distributed Computing*, pages 377–408, 1990.
- [11] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer-Verlag New York, Inc., New York, NY, USA, 1992.
- [12] A. Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, SFCS ’77*, pages 46–57, Washington, DC, USA, 1977. IEEE Computer Society.
- [13] M. Rungger and P. Tabuada. Abstracting and refining robustness for cyber-physical systems. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control, HSCC ’14*, pages 223–232, New York, NY, USA, 2014. ACM.
- [14] D. C. Tarraf, A. Megretski, and M. A. Dahleh. A framework for robust stability of systems over finite alphabets. *IEEE Transactions on Automatic Control*, 53(5):1133–1146, June 2008.
- [15] K. W. Wong, R. Ehlers, and H. Kress-Gazit. Correct high-level robot behavior in environments with unexpected events. In *Robotics: Science and Systems Conference (RSS’14)*, 2014.
- [16] K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.

Appendices

A. PROOF OF PROPOSITION 5

PROOF. Let $\sigma \in \Sigma^\omega$ and $m \in M^\omega$ be such that $(\sigma_k, m_k) = \bar{\sigma}_k$. When $\bar{\sigma} \models \square \diamond \bar{1}_{jump}$, the proposition holds trivially. Consider the case when $\bar{\sigma} \models \square \neg \bar{1}_{jump}$. This means that there are only finite instances of $\bar{\sigma}_k \bar{\sigma}_{k+1} \models \bar{1}_{jump}$ i.e in other words, only for finite $k > 0$, $(\sigma_k, m_k) \in \text{ExtTs}(\sigma_{k-1}, m_{k-1})$. Let there be $r < \infty$ such instances i.e $\exists i_l \in \mathbb{N}$ such that $\bar{\sigma}_{i_{l-1}} \bar{\sigma}_{i_l} \models \bar{1}_{jump}$ for $l \in \{1, 2, \dots, r\}$.

Without loss of generality, assume $0 < i_1 < i_2 < \dots < i_r$. $\forall l \in \{1, 2, \dots, r-1\} \exists q_l \in \{0, 1, 2, \dots, n\}$ such that $(\sigma_i \cap \text{AP}_{\text{sys}}, m_i) = g_{q_l}(\sigma_i \cap \text{AP}_{\text{env}}, m_{i-1})$ for $i \in \{i_j, i_{j+1}, \dots, i_{j+1} - 2\}$. Otherwise at some i between i_j and $i_{j+1} - 2$, $\bar{\sigma}_i \bar{\sigma}_{i+1} \models \bar{1}_{jump}$ - contradicting our assumption. Similarly, $\forall i \in \{0, 1, \dots, i_1 - 2\}$, $(\sigma_i \cap \text{AP}_{\text{sys}}, m_i) = g_0(\sigma_i \cap \text{AP}_{\text{env}}, m_{i-1})$.

$\forall l \in \{1, 2, \dots, r-1\}$, $(\sigma_{i_l}, m_{i_l}) \in \text{ExtTs}(\sigma_{i_{l-1}}, m_{i_{l-1}})$ by the definition of $\bar{1}_{jump}$. From the definition of ExtTs we can conclude that $(\sigma_{i_l}, m_{i_l}) \in I_{\varphi_{q_l}}$ where $m_{i_l} \in M_{q_l}$ i.e σ_{i_l} is φ_{q_l} -reachable. For $l \in \{1, 2, \dots, r\}$ by definition of φ_i -reachability and the discussion above, there exists $\beta^l \in \text{Plays}(g_{q_l})$ and k_l such that $\sigma_{i_l} = \beta^l_{k_l}$, $\beta^l \models \Theta_0^{\text{env}}$, and

$$\beta^l_{j:(j+1)} \models \rho^{\text{env}} \text{ for } j < k-1. \quad (21)$$

and $\forall l \leq r-1$

$$\beta^l_{k_l+w} = \sigma_{i_l+w} \text{ for } 0 \leq w < i_{l+1} - i_l \quad (22)$$

For $l = r$

$$\beta^r_{k_r+w} = \sigma_{i_r+w} \text{ for } 0 \leq w \quad (23)$$

since $\bar{\sigma}_{i_r} \models \neg \square \bar{1}_{jump}$.

We write $\tau^l \xi^l \alpha^l = \beta^l$ by taking $\tau_j^l = \beta_j^l$ for $0 \leq j < k$. For $l \leq r-1$, $\xi_w^l = \beta^l_{k_l+w}$ where $0 \leq w < i_{l+1} - i_l$, and $\alpha^l_{j-k_l-i_{l+1}+i_l} = \beta_j^l$ for $j \geq k_l + i_{l+1} - i_l$. For $l = r$, $(\xi^r \alpha^r)_{j-k_r} = \beta_j^r$ for $j \geq k_r$. To conclude the setup, we argued that $\sigma = \tau^0 \xi^1 \xi^2 \dots \xi^r \alpha^r$ where $\tau^0 \in \text{Pref}(g_0)$. Combine with m as in constructing $\bar{\sigma}$ to have $\bar{\sigma} = \bar{\tau}^0 \bar{\xi}^1 \bar{\xi}^2 \dots \bar{\xi}^n \bar{\alpha}^r$

We want to show that

$$\begin{aligned} & \bar{\tau}^0 \bar{\xi}^1 \bar{\xi}^2 \dots \bar{\xi}^k \bar{\xi}^{k+1} \dots \bar{\xi}^n \bar{\alpha}^n \models \\ & \Theta_0^{\text{env}} \square (\rho^{\text{env}} \vee \bar{1}_{jump}) \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right) \\ & \implies \Theta_0^{\text{sys}} \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right) \quad (24) \end{aligned}$$

This is equivalent to at least one of the following subformulae being satisfied: $\neg \Theta_0^{\text{env}}$, $\diamond (\neg \rho^{\text{env}} \wedge \neg \bar{1}_{jump})$,

$\diamond \square \neg \psi_j^{\text{env}}$ for some j , or

$$\Theta_0^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right). \quad (25)$$

Since $\tau^0 \in \text{Pref}(g_0)$ by hypothesis, there exists $\gamma \in \Sigma^\omega$ such that $\tau^0 \gamma \in \text{Plays}(g_0)$. Also by hypothesis, (g_0, m_0) realizes φ_0 , i.e., $\text{Plays}(g_0) \subseteq L(\varphi_0)$, hence $\tau^0 \gamma \models \varphi_0$. Note that $|\tau^0| \geq 1$ since $(\tau^0 \cap \text{AP}_{\text{env}}, m_1) = g_0(\tau^1 \cap \text{AP}_{\text{sys}}, m_0)$ by the definition of Plays f . $\tau^0 \gamma \models \neg \Theta_0^{\text{env}}$ if and only if $\bar{\sigma} \models \neg \Theta_0^{\text{env}}$. Thus, if $\tau^0 \gamma \models \neg \Theta_0^{\text{env}}$, then (24) holds directly. Otherwise (i.e., if $\tau^0 \gamma \models \Theta_0^{\text{env}}$), consider the subformula $\diamond (\neg \rho^{\text{env}} \wedge \neg \bar{1}_{jump})$.

For all $k < |\tau^0| - 1$, $\bar{\sigma}_{k:k+1} \models \neg \bar{1}_{jump}$ by construction. Also, for $k \geq \bar{k} = |\tau^0 \xi^1 \xi^2 \dots \xi^k \xi^{k+1} \dots \xi^r|$, $\bar{\alpha}_{k-\bar{k}, k+1-\bar{k}}^r \models \neg \bar{1}_{jump}$ and $\bar{\xi}_{-1}^r \bar{\alpha}_0^r \models \neg \bar{1}_{jump}$. From the decomposition given above, $\bar{\tau}_{-1}^0 \bar{\xi}_0^1 \models \bar{1}_{jump}$ and $\bar{\xi}_{-1}^l \bar{\xi}_0^{l+1} \models \bar{1}_{jump}$ for $l \leq r-1$. For this case, we can conclude

$$\begin{aligned} \bar{\sigma} \models \diamond (\neg \rho^{\text{env}} \wedge \neg \bar{1}_{jump}) & \implies (\tau^0 \gamma \models \diamond \neg \rho^{\text{env}}) \vee \\ & (\xi^1 \xi^2 \dots \xi^k \xi^{k+1} \dots \xi^r \alpha^r \models \diamond \neg \rho^{\text{env}}) \end{aligned}$$

If $\tau^0 \gamma \models \diamond \neg \rho^{\text{env}}$, then there is a minimum k such that $\sigma_{k:\gamma} \models \neg \rho^{\text{env}}$ or $\gamma_{(k-|\sigma|):} \models \neg \rho^{\text{env}}$. If $k < |\sigma| - 1$, then $\bar{\sigma} \models \diamond (\neg \rho^{\text{env}} \wedge \neg \bar{1}_{jump})$ (recall $\bar{\sigma}_{k:k+1} \models \neg \bar{1}_{jump}$). Then (24) is satisfied.

From the definition of ExtTs and $\bar{1}_{jump}$ we have

$$\tau_{-1}^0 \xi_0^1 \models \rho^{\text{sys}}, \quad (26)$$

and

$$\xi_{-1}^l \xi_0^{l+1} \models \rho^{\text{sys}} \forall l \in \{1, 2, 3, \dots, r-1\}, \quad (27)$$

which we will refer to later while addressing the final case. The other case in which $\bar{\sigma} \models \diamond (\neg \rho^{\text{env}} \wedge \neg \bar{1}_{jump})$ is if $\xi^l \models \diamond \neg \rho^{\text{env}}$ for some l or $\xi^r \alpha^r \models \diamond \neg \rho^{\text{env}}$ (recall $\bar{\tau}_{-1}^0 \xi_0^1 \models \bar{1}_{jump}$ and $\xi_{-1}^l \xi_0^{l+1} \models \bar{1}_{jump}$ for $l \leq r-1$). If for any $d < |\xi^l| - 1$ for $l < r$, $\xi_d^l \xi_{d+1}^l \models \neg \rho^{\text{env}}$, then (24) is directly satisfied. If $\xi^n \alpha^n \models \diamond \neg \rho^{\text{env}}$ then again (24) is satisfied.

Otherwise, suppose that $\bar{\sigma} \models \diamond \square \neg \psi_j^{\text{env}}$ for some j . From the semantics of LTL and the fact that ψ_j^{env} contains no temporal operators, this implies $\alpha^n \models \diamond \square \neg \psi_j^{\text{env}} \leftrightarrow \bar{\sigma} \models \diamond \square \neg \psi_j^{\text{env}}$. In this case where $\alpha^n \models \diamond \square \neg \psi_j^{\text{env}}$, (24) again holds.

We now consider the final case where $\bar{\sigma} \models \Theta_0^{\text{env}} \wedge \square \diamond (\rho^{\text{env}} \vee \bar{1}_{jump}) \wedge \left(\bigwedge_{j=1}^J \square \diamond \psi_j^{\text{env}} \right)$. By φ_{q_r} -reachability, $\tau_{d,d+1}^r \models \rho^{\text{env}}$ for all $d < |\tau^r| - 1$. And, $\xi^r \alpha^r \models \square \rho^{\text{env}}$ as argued for this case (otherwise (24) would directly hold). Thus, $\tau^r \xi^r \alpha^r \models \square \rho^{\text{env}}$. Recall that $\tau^r \xi^r \alpha^r \models \Theta_{q_l}^{\text{env}}$. Because $\tau^r \xi^r \alpha^r \in \text{Plays}(f_{q_r})$ and $\tau^r \xi^r \alpha^r \models \Theta_{q_l}^{\text{env}}$, if neither $\tau^r \xi^r \alpha^r \models \diamond \neg \rho^{\text{env}}$ nor $\diamond \square \neg \psi_j^{\text{env}}$ for any j , it must be that $\tau^r \xi^r \alpha^r$ satisfies $\Theta_{q_r}^{\text{sys}} \wedge \square \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \square \diamond \psi_k^{\text{sys}} \right)$.

By φ_{q_l} -reachability of ξ_0^l and strict-realizability, $\xi_{d,d+1}^l \models \rho^{\text{sys}} \forall d < |\xi^l| - 1$ for $l \leq r - 1$. From (26) and (27), it follows that

$$\tau_{-1}^0 \xi^1 \xi^2 \dots \xi^n \alpha \models \Box \rho^{\text{sys}} \wedge \left(\bigwedge_{k=1}^K \Box \Diamond \psi_k^{\text{sys}} \right).$$

Recall the suffix γ such that $\tau^0 \gamma \in \text{Plays}(g_0)$. (g_0, m_0) strictly realizes φ_0 , therefore if $\tau^0 \gamma \models \Theta_0^{\text{env}}$, it must be that $\tau^0 \gamma \models \Theta_0^{\text{sys}}$. $\bar{\sigma} \models \Theta_0^{\text{sys}}$ since $\bar{\sigma} \models \Theta_0^{\text{env}}$ (9). Furthermore, because in this case we are assuming there is no $0 \leq k < |\tau^0| - 1$ such that $\tau_{k:(k+1)}^0 \models \neg \rho^{\text{env}}$ (otherwise we would have (24) hold directly), it follows from strict realizability (cf. (10)) that for $0 \leq k < |\tau^0| - 1$, $\tau_{k:(k+1)}^0 \models \rho^{\text{sys}}$, and therefore the condition in the proposition for $\bar{\sigma}$ holds.

□

B. PROOF OF PROPOSITION 6

PROOF. Let $\sigma \in \Sigma^\omega$ and $m \in M^\omega$ be such that $(\sigma_k, m_k) = \bar{\sigma}_k$. Repeating the arguments in Section 4.1, we only need to reason about the case where $\bar{\sigma} \models \Diamond \Box \neg \bar{1}_{\text{jump}}$. In this case, as before, let there be $r < \infty$ instances in which $\bar{1}_{\text{jump}}$ holds. Let these instances be i_l for $l \leq r$ such that $0 < i_1 < i_2 < \dots < i_r$. By definition, $\bar{\sigma}_{i_l-1} \bar{\sigma}_{i_l} \models \bar{1}_{\text{jump}} \implies \bar{\sigma}_{i_l} \in \text{ExtTs}_{\text{reach}}(\bar{\sigma}_{i_l-1})$. Reasoning as before, for $l \in \{1, 2, \dots, r\}$, $\exists q_l$ such that $\beta^l \in \text{Plays}(h_{q_l})$ and k_l such that $\sigma_{i_l} = \beta_{k_l}^l, \beta^l \models \Theta_0^{\text{env}}$ and conditions (21),(22),(23) hold.

As before, we write $\tau^l \xi^l \alpha^l = \beta^l$. Then, repeating the arguments from earlier decompose $\bar{\sigma}$ as $\bar{\sigma} = \bar{\tau}^0 \bar{\xi}^1 \dots \bar{\xi}^r \bar{\alpha}^r$ where $\tau^0 \in \text{Pref}(h^0)$. Note that the decomposition was such that for $l \in \{1, 2, \dots, r\}$, $\xi_0^l \in I_{\varphi_{\text{reach}}^{q_l}}^{M-r}$ holds if $q_l \geq 1$ and $\xi^l \in I_\varphi^M$ if $q_l = 0$. That is to say each ξ_0^l is $\varphi_{\text{reach}}^{q_l}$ -reachable or φ -reachable, as the case may be. We want to show (24).

Reasoning about the other trivial cases as before and rejecting them, consider the case when

$$\bar{\sigma} \models \Theta_0^{\text{env}} \Box (\rho^{\text{env}} \vee \bar{1}_{\text{jump}}) \wedge \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$$

$\bar{\sigma} \models \Theta_0^{\text{env}} \leftrightarrow \bar{\tau}^0 \models \Theta_0^{\text{env}}$. $\tau^0 \in \text{Pref}(h^0) \rightarrow \exists \gamma \in \Sigma^\omega$. $\tau^0 \gamma \in \text{Plays}(h^0)$. By strict realizability, we have $\tau^0 \gamma \models \Theta_0^{\text{env}} \rightarrow \tau \gamma \models \Theta_0^{\text{sys}}$. Therefore, $\bar{\sigma} \models \Theta_0^{\text{sys}}$.

For $2 \leq l \leq r - 1$, we have $\xi_0^l \xi_0^{l+1} \models \rho^{\text{sys}}$ from the definition of $\bar{1}_{\text{jump}}$ and $\text{ExtTs}_{\text{reach}}$. We also have $\tau_{-1}^0 \xi_0^1 \models \rho^{\text{sys}}$. When $\bar{\sigma} \models \Box (\rho^{\text{env}} \vee \bar{1}_{\text{jump}})$, by strict realizability and $\varphi_{\text{reach}}^{q_l}$ -realizability (or φ -reachability as the case maybe) for $l \leq r - 1$, we have $\xi_{d,d+1}^l \models \rho^{\text{sys}} \forall d < |\xi^l| - 1$ since $\xi_{d,d+1}^l \models \rho^{\text{env}} \forall d < |\xi^l| - 1$. Also, we have $\tau_{d,d+1}^0 \models \rho^{\text{sys}}$ for $\forall d < |\tau^0| - 1$ since $\tau_{d,d+1}^0 \models \rho^{\text{env}} \forall d < |\tau^0| - 1$.

Consider $\tau^r \xi^r \alpha^r \in \text{Plays}(h^{q_r})$. If $q_r > 0$, h^{q_r} strictly realizes $\varphi_{\text{reach}}^{q_r}$ and ξ_0^r is $\varphi_{\text{reach}}^{q_r}$ -reachable. If $q_r = 0$, h^0 strictly-realizes φ and ξ_0^r is φ -reachable. Therefore, if $\xi^r \alpha^r \models \Box \rho^{\text{env}}$ we have $\tau^r \xi^r \alpha^r \models \Box \rho^{\text{env}}$. By strict realizability, we have $\tau^r \xi^r \alpha^r \models \Box \rho^{\text{sys}}$.

To conclude, we showed $\bar{\sigma} \models \Theta_0^{\text{env}} \wedge \Box (\rho^{\text{env}} \vee \bar{1}_{\text{jump}}) \rightarrow \bar{\sigma} \models \Theta_0^{\text{sys}} \wedge \Box \rho^{\text{sys}}$. Now only showing liveness remains. If $q_r > 0$, we have $\tau^r \xi^r \alpha^r \models \Diamond \mathcal{T}_{\text{reach}}$ or $\tau^r \xi^r \alpha^r$ satisfies $\neg \Theta_0^{\text{env}}$ or $\Diamond \neg \rho^{\text{env}}$ or $\Diamond \Box \neg \psi_j^{\text{env}}$ for some j since h^{q_r} realizes $\varphi_{\text{reach}}^{q_r}$. But, as an assumption we have $\bar{\sigma} \models \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$ and by the semantics of LTL we have $\alpha^r \models \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$. And, $\alpha^r \models \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right) \leftrightarrow \tau^r \xi^r \alpha^r \models \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$. By $\varphi_{\text{reach}}^{q_r}$ -reachability of ξ_0^r , we have $\tau^r \xi^r \alpha^r \models \Theta_{q_r}^{\text{env}}$. And, above we reasoned $\tau^r \xi^r \alpha^r \models \Box \rho^{\text{env}}$ for this case. Therefore, it must be the case that $\tau^r \xi^r \alpha^r \models \Diamond \mathcal{T}_{\text{reach}}$. Observe that if $s \in I_\varphi(h^0, m_0^0)$ then $\exists w \in M^0. (s, w) \in I_\varphi^M(h^0, w_0^0)$. $\tau^r \xi^r \alpha^r \models \Diamond \mathcal{T}_{\text{reach}}$ leads to a contradiction as by the definition of $\text{Plays}(\bar{h})$, we should switch strategies when we reach a state s that satisfies $\mathcal{T}_{\text{reach}}$. Therefore, $q_r = 0$ when $\bar{\sigma} \models \Theta_0^{\text{env}} \wedge \Box (\rho^{\text{env}} \vee \bar{1}_{\text{jump}}) \wedge \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$.

By strict realizability and φ -reachability, we reasoned that $\tau^r \xi^r \alpha^r \models \Theta_0^{\text{env}} \wedge \Box \rho^{\text{env}} \rightarrow \bar{\sigma} \models \Theta_0^{\text{sys}} \wedge \Box \rho^{\text{sys}}$. Additionally $\tau^r \xi^r \gamma^r \models \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{env}} \right)$, by which we have $\tau^r \xi^r \alpha^r \models \Theta_0^{\text{sys}} \wedge \Box \rho^{\text{sys}} \wedge \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{sys}} \right)$ (since $q_r = 0$ and from the definition of (h^0, m_0^0) realizing φ). Therefore, we have $\bar{\sigma} \models \Theta_0^{\text{sys}} \wedge \Box \rho^{\text{sys}} \wedge \left(\bigwedge_{j=1}^J \Box \Diamond \psi_j^{\text{sys}} \right)$, hence proving the proposition. □