

Interpolation-Based GR(1) Assumptions Refinement

Davide G. Cavezza^(✉) and Dalal Alrajeh

Imperial College London, London, UK
{d.cavezza15,dalal.alrajeh}@imperial.ac.uk

Abstract. This paper considers the problem of assumptions refinement in the context of unrealizable specifications for reactive systems. We propose a new counterstrategy-guided synthesis approach for GR(1) specifications based on Craig’s interpolants. Our interpolation-based method identifies causes for unrealizability and computes assumptions that directly target unrealizable cores, without the need for user input. Thereby, we discuss how this property reduces the maximum number of steps needed to converge to realizability compared with other techniques. We describe properties of interpolants that yield helpful GR(1) assumptions and prove the soundness of the results. Finally, we demonstrate that our approach yields weaker assumptions than baseline techniques.

Keywords: Reactive synthesis · Assumption refinement · Interpolation

1 Introduction

Constructing formal specifications that capture user requirements precisely and from which implementations can be successfully derived is a difficult task [25]. Their imprecision often results from the conception of over-ideal systems, i.e., where the environment always behaves as expected [2, 26]. Thus one of the challenges in building correct specifications is identifying sufficient assumptions over the environment under which a system would always be able to guarantee their satisfaction, in other words making a specification *realizable*.

This paper presents a new technique for automatically synthesizing assumptions over an adversarial environment for realizability assurance. More specifically, we develop a novel counterstrategy-guided synthesis procedure that iteratively generates assumption refinements, expressed in a fragment of Linear Temporal Logic (LTL) called *Generalized Reactivity (1)* (GR(1) for short), based on logical interpolation. *Craig interpolants* characterize automatically computable explanations for the inconsistency between Boolean formulae, in their shared alphabet. We exploit this feature to construct expressions that *explain* why a counterstrategy, and hence the environment, violates a guarantee, and whose negations form assumptions.

We demonstrate in our case study applications that our approach *directly targets unrealizable cores*, in the sense that by adding the assumptions returned at

each iteration, a specific subset of *minimally unfulfillable guarantees* [13] becomes realizable. Therefore each iteration takes a step closer to realizability. To characterize the scope of our approach we introduce the notion of fully-separable interpolants and prove the soundness of our computation when interpolants are fully separable. We further provide a discussion about the complexity of the proposed approach and its convergence, as well as the weakness of our refinements in comparison with those computed by existing techniques [3, 4, 27].

2 Related Work

Recent years have seen the development of effective counterstrategy-guided approaches to GR(1) assumptions refinement, notably [3, 4, 27]. Nonetheless those approaches depend significantly on users' knowledge of the problem domain and of the cause of unrealizability. The work in [27] requires users to specify a set of temporal logic templates as formulae with placeholders to be replaced with Boolean variables. Assumptions are then generated as instantiations of such templates that eliminate a given counterstrategy. This typically constrains the search space to only a subset of GR(1) formulae, which do not necessarily address the cause of unrealizability, and potentially eliminate viable solutions to the realizability problem. Similarly, the work in [3], although generating such templates automatically, requires users to provide a subset of variables for template instantiation. Unless the user knows the exact subset of variables that form the cause, this may yield assumptions that do not target the true cause of unrealizability, leading to refinements that needlessly over-constrain the environment. Our proposed method instead directly targets counter-strategies and unrealizable cores, and does not require users to provide variables for constructing refinements.

Other related work on assumption refinement includes those operating directly on game structures [12]. With regard to the parity game model used for controller synthesis (such as in [32, 33]), the paper defines the concept of safety assumptions as sets of edges that have to be avoided by the environment, and the concept of fairness assumptions as sets of edges that have to be traversed by the environment infinitely often. The work devises an algorithm for finding minimal edge sets in order to ensure that the controller has a winning strategy. Our approach instead focuses on synthesizing general declarative temporal assertions whose inclusion has the effect of removing edges from the game structure, and directly targeting sources of unrealizability. The problem of synthesizing environment constraints has been tackled in the context of assume-guarantee reasoning for compositional model checking [15, 21, 31] to support compositional verification. In these, assumptions are typically expressed as LTSs and learning algorithms like \mathcal{L}^* [5] are used to incrementally refine the environment assumptions needed in order to verify the satisfaction of properties.

Craig interpolants have been deployed in the context of abstraction refinement for verification in [18, 19]. The differences with our work are in specification language and overall objective: they seek additional assertions for static analysis of programs, while we look for GR(1) refinements of systems specifications

to enable their automated synthesis. The authors of [17] use interpolation to support the extraction of pre- and trigger-conditions of operations within event-driven systems to enable the ‘satisfaction’ of goals expressed within restricted fragment of LTL. Though different in objective, approach and class of properties, our technique can help in identifying specifications operationalizable by [17].

3 Background

Generalized Reactivity (1) Specifications. LTL [29] is a formalism widely used for specifying reactive systems. The syntax of LTL is defined over a finite non-empty set of propositional variables \mathcal{V} , the logical constants *true* and *false*, Boolean connectives, and operators **X** (next), **G** (always), **F** (eventually), **U** (until). Given a set of states Q and a labelling function $\lambda : Q \rightarrow 2^{\mathcal{V}}$, an LTL formula ϕ is interpreted over an infinite sequence of states $\sigma = q_0q_1\dots$ in the standard way, and its language $\mathcal{L}(\phi)$ is the set of (infinite) state valuation sequences $w = \lambda(q_0)\lambda(q_1)\dots$ such that $w \models \phi$. We assume that the set \mathcal{V} consists of two disjoint sets: *input variables* \mathcal{X} and *output variables* \mathcal{Y} . We will use the expression $B(\mathcal{V})$ for a boolean expression (i.e., a logical expression without temporal operators) which uses variables in the set \mathcal{V} . We will also denote by $\mathbf{X}\mathcal{V}$ the set of expressions obtained by prepending a “next” operator to the variables in \mathcal{V} : this is equivalent to the set of primed versions of such variables [9].

Generalized Reactivity (1) specifications (written GR(1) for short) are a subset of LTL of the form $\phi^{\mathcal{E}} \rightarrow \phi^{\mathcal{S}}$ where $\phi^{\mathcal{E}}$ represents the assumptions of an environment and $\phi^{\mathcal{S}}$ the guarantees of a controller. The expression ϕ^{θ} , where $\theta \in \{\mathcal{E}, \mathcal{S}\}$, is specified as conjunction of the following: (1) a Boolean formula φ_{init}^{θ} of the form $B(\mathcal{X})$ if $\theta = \mathcal{E}$ and $B(\mathcal{V})$ otherwise, representing initial conditions; (2) a set of LTL formulae φ_{inv}^{θ} of the form $\mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{X})$ if $\theta = \mathcal{E}$ and $\mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{Y})$ when $\theta = \mathcal{S}$, representing invariants; and (3) a set of LTL formulae φ_{fair}^{θ} of the form $\mathbf{G}FB(\mathcal{V})$ representing fairness conditions. We will sometimes indicate GR(1) specifications as a tuple $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ with $\phi^{\theta} = \{\varphi_{init}^{\theta}\} \cup \{\varphi_{inv}^{\theta}\} \cup \{\varphi_{fair}^{\theta}\}$.

A finite-state Moore transducer is a tuple $M = \langle Q, q_0, \mathcal{I}, \mathcal{O}, \rho, \delta \rangle$ where Q is a set of states, $q_0 \in Q$ is the initial state, $\rho : Q \times \mathcal{I} \rightarrow Q$ is the transition function, and $\delta : Q \rightarrow \mathcal{O}$ is the output function. Given an input sequence $w = i_0i_1\dots$, a run of M is the sequence $\sigma = q_0q_1\dots$ such that $q_{k+1} = \rho(q_k, i_k)$ for all $k \geq 0$. A run σ on input sequence $w \in \mathcal{I}^{\omega}$ produces an infinite word $M(w) = (\delta(q_0), i_0), (\delta(q_1), i_1)\dots$. The language of a Moore transducer M is $\mathcal{L}(M) = \{M(w) | w \in \mathcal{I}^{\omega}\}$, i.e., the infinite words generated by a sequence of inputs and the corresponding outputs over runs of M . A Moore transducer M is said to satisfy an LTL expression ϕ if $\mathcal{L}(M) \subseteq \mathcal{L}(\phi)$; in this case we also say that M is a model of ϕ and we denote it as $M \models \phi$. A GR(1) property ϕ is said to be realizable if there exists an M (representing a controller) such that $M \models \phi$.

Given a specification $\langle \phi^{\mathcal{E}}, \phi^{\mathcal{S}} \rangle$ that is unrealizable, we say that $\varphi^{\mathcal{S}} \subseteq \phi^{\mathcal{S}}$ is *minimally unfulfillable w.r.t. to $\phi^{\mathcal{E}}$* iff the removal of any guarantee $g \in \varphi^{\mathcal{S}}$ makes $\langle \phi^{\mathcal{E}}, \varphi^{\mathcal{S}} \setminus \{g\} \rangle$ realizable [13]. Furthermore, an assumption $a \in \phi^{\mathcal{E}}$ is said to be *unhelpful w.r.t. $\phi^{\mathcal{S}}$* if $\forall \varphi^{\mathcal{S}} \subseteq \phi^{\mathcal{S}}. \langle \phi^{\mathcal{E}}, \varphi^{\mathcal{S}} \rangle$ is realizable $\leftrightarrow \langle \phi^{\mathcal{E}} \setminus \{a\}, \varphi^{\mathcal{S}} \rangle$ is

realizable. It is said to be *helpful* otherwise. Given a set of minimally unfulfillable guarantees φ^S w.r.t. ϕ^E , let $\varphi^E \subseteq \phi^E$ be a set of helpful assumptions for φ^S ; the specification $\langle \varphi^E, \varphi^S \rangle$ is called an *unrealizable core* [13].

If a specification ϕ over $\mathcal{V} = \mathcal{X} \cup \mathcal{Y}$ is unrealizable, an unrealizable core $\langle \varphi^E, \varphi^S \rangle$ and an environment strategy (called a *counterstrategy*) can be computed [13,23]. A *counterstrategy* is defined as a Moore transducer $(S, s_{init}, 2^{\mathcal{Y}'}, 2^{\mathcal{X}}, \rho, \delta)$ that satisfies φ^E and violates φ^S [4]. It describes the inputs produced by an admissible environment in response to the output configuration yielded by the controller in order to force the violation of ϕ . The runs of a counterstrategy are called *plays*. The terms ‘counterstrategy’ and ‘play’ come from the game-theoretic algorithms used to reason about realizability [3,9,23]. The transition function ρ depends only on a subset of the output variables $\mathcal{Y}' \subseteq \mathcal{Y}$ [23]. We define a labelling function $\lambda' : S \rightarrow 2^{\mathcal{X} \cup \mathcal{Y}'}$ over states in the counterstrategy in this way: a propositional variable is in $\lambda'(s)$ if it is asserted in all the incoming transitions of s , while $\lambda'(s)$ is arbitrary for any s with no incoming transitions.

Interpolants. Craig interpolation was originally defined for first-order logic [16] and later for propositional logic [24]. No interpolation theorems have been proved for the general LTL. Extensions have been proposed recently for LTL fragments [20,22]. However these do not include GR(1) formulae and therefore are not applicable in our case. We use interpolation for propositional logic.

Formally, given an unsatisfiable conjunction of formulae $\alpha \wedge \beta$, a Craig interpolant I is a formula that is implied by α , is unsatisfiable in conjunction with β , and is defined on the common alphabet of α and β . We write \mathcal{L}_ϕ to denote the set of variables that occur in a formula ϕ (also called the *alphabet* of ϕ).

Definition 1 (Interpolant [24]). *Let α and β be two logical formulae such that their conjunction $\alpha \wedge \beta$ is unsatisfiable. Then there exists a third formula I , called interpolant of α and β , such that, $\alpha \rightarrow I$, $I \rightarrow \neg\beta$ and $\mathcal{L}_I \subseteq \mathcal{L}_\alpha \cap \mathcal{L}_\beta$.*

An interpolant can be considered as an over-approximation of α that is still unsatisfiable in conjunction with β . As stated in Craig’s interpolation theorem, although an interpolant always exists, it is not unique. Several efficient algorithms have been proposed for interpolation in propositional logics. The resulting interpolant depends on the internal strategies of these algorithms (e.g., SAT solvers, theorem provers). Our approach is based on McMillan’s interpolation algorithm described in [30] and implemented in MathSAT [14]. In brief, the algorithm considers a proof by resolution for the unsatisfiability of $\alpha \wedge \beta$.

4 Approach Overview

The general procedure is based on a sequence of realizability checks and counterstrategy computations, in the spirit of [3,27]. A specification $\langle \phi^E, \phi^S \rangle$ is first checked for realizability. If it is unrealizable, a counterstrategy C and an unrealizable core $\langle \varphi^E, \varphi^S \rangle$ are computed. The counterstrategy constitutes an example

of environment behaviours that force the violation of the guarantees of φ^S : therefore, the assumptions φ^E are refined by adding a GR(1) formula which is inconsistent with the counterstrategy. A set of such formulae Ψ is automatically computed by interpolating (α) the description of an environment behaviour in the counterstrategy, given by the assumptions and a sequence of state labellings in the counterstrategy; and (β) the guarantees, and by negating the interpolant. A formula $\psi_i \in \Psi$ is added to the original set of assumptions ϕ^E and the procedure repeats the above steps recursively until realizability is achieved. Algorithm 1 describes this procedure schematically.

Algorithm 1. CounterstrategyGuidedRefinement procedure

Data: ϕ^E , assumptions
Data: ϕ^S , guarantees
Result: $\{\psi_i\}$, set of alternative assumption refinements such that $\phi^E \wedge \psi_i \rightarrow \phi^S$ is realizable for every i

```

1 if Satisfiable( $\phi^E \rightarrow \phi^S$ )  $\&$  not Realizable( $\phi^E \rightarrow \phi^S$ ) then
2   | ( $\varphi^E, \varphi^S, C$ ) := Counterstrategy( $\phi^E, \phi^S$ );
3   |  $\Psi$  := InterpolationBasedSynthesis( $\varphi^E, \varphi^S, C$ );
4   | foreach  $\psi_i \in \Psi$  do
5     |   | foreach  $\psi'_j \in$  CounterstrategyGuidedRefinement( $\phi^E \wedge \psi_i, \phi^S$ ) do
6       |   |   | refinements.add( $\psi_i \wedge \psi'_j$ );
7     |   |   | end
8     |   | end
9     |   | return refinements;
10 else if Satisfiable( $\phi^E \rightarrow \phi^S$ )  $\&$  Realizable( $\phi^E \rightarrow \phi^S$ ) then
11 |   | return  $\{true\}$ ;
12 else
13 |   | return  $\{false\}$ ;

```

The function **InterpolationBasedSynthesis** constitutes the core of our proposal (see Algorithm 2). It takes as inputs an unrealizable core and a counterstrategy and executes the computation of Ψ via interpolation. We give the details in the following section.

5 Interpolation-Based Synthesis

Each execution of **InterpolationBasedSynthesis** involves extracting temporal formulae that are satisfied by a single play of a counterstrategy (henceforth called *counterplay*), and obtaining refinements from its negation. It is sufficient to exclude a single counterplay of a counterstrategy to eliminate the entire counterstrategy from models of the assumption. Reasoning about counterplays has also some advantages, which are discussed in Sect. 8. For the purpose of this paper, we assume that the procedure **ExtractCounterplay** (line 1) extracts a counterplay π_C at random and consider metrics for selecting one in future work. A counterplay representing the violation of an initial condition or an invariant

Algorithm 2. $\text{InterpolationBasedSynthesis}(\varphi^{\mathcal{E}}, \varphi^{\mathcal{S}}, C)$

Data: $\varphi^{\mathcal{E}}$, environment assumptions (in an unrealizable core)
Data: $\varphi^{\mathcal{S}}$, system guarantees (in an unrealizable core)
Data: C , counterstrategy
Result: Ψ , alternative assumptions eliminating the counterstrategy

```

1  $\pi_C := \text{ExtractCounterplay}(C);$ 
2  $u := 0;$ 
3  $\pi_{C,u} := \pi_C;$ 
4  $\Psi_{old} := \emptyset;$ 
5  $\text{stopping\_condition} := \text{true};$ 
6 repeat
7    $[[\pi_{C,u}, \varphi_u^{\mathcal{E}}]] := \text{TranslateCounterplayAssumptions}(\pi_{C,u}, \varphi^{\mathcal{E}});$ 
8    $[[\varphi_u^{\mathcal{S}}]] := \text{TranslateGuarantees}(\pi_{C,u}, \varphi^{\mathcal{S}});$ 
9    $I_u := \text{Interpolate}([[ \pi_{C,u}, \varphi_u^{\mathcal{E}} ]], [[ \varphi_u^{\mathcal{S}} ]]);$ 
10  if  $I_u == \text{false}$  or  $I_u$  is not fully-separable then
11     $\Psi := \{\text{false}\};$ 
12     $\text{stopping\_condition} := \text{true};$ 
13  else
14     $\mathcal{T}(I_u) := \text{TranslateInterpolant}(\pi_{C,u}, I_u);$ 
15     $\Psi := \text{ExtractDisjuncts}(\neg \mathcal{T}(I_u));$ 
16    if  $\pi_{C,u}$  is looping then
17      if  $\Psi \neq \Psi_{old}$  then
18         $\Psi_{old} := \Psi;$ 
19         $u := u + 1;$ 
20         $\pi_{C,u} := \text{UnrollCounterplay}(\pi_C, u);$ 
21         $\text{stopping\_condition} := \text{false};$ 
22      else
23         $\text{stopping\_condition} := \text{true};$ 
24      end
25    end
26  end
27 until  $\text{stopping\_condition};$ 
28 return  $\Psi;$ 

```

guarantee is finite, while that of a fairness guarantee violation ends in a loop [28]. We call the latter a *looping counterplay*, and the loop an *ending loop*.

We distinguish four types of states that may appear in π_C : (a) the *initial state* $S^{init} = \{s^{init}\}$; (b) the *failing state* in a finite counterplay $S^{fail} = \{s^{fail}\}$ (c) *looping states* that include the states in ending loop, $S^{loop} = \{s_1^{loop}, \dots, s_h^{loop}\}$, (d) *transient states* including all states between the initial state and the first failing state or loop state (exclusive) $S^{trans} = \{s_1^{trans}, \dots, s_k^{trans}\}$. With this classification, a finite counterplay has the form $s^{init} s_1^{trans} \dots s_k^{trans} s^{fail}$; whilst a looping counterplay has the form $s^{init} s_1^{trans} \dots s_k^{trans} (s_1^{loop} \dots s_h^{loop})^\omega$. The formulae in the next subsection also refer to a fifth set of states, called *unrolled states*, which represent replicates of looping states, and to the unrolling degree u . They are explored in Sect. 5.2. Each state in π_C is labelled with variables from

the set $\mathcal{X} \cup \mathcal{Y}'$ defined in Sect. 3. The value of u is initialized to 0, and thus $\pi_{C,u}$ equates to π_C (lines 2–3).

The extraction of the counterplay occurs at the start of every call of the synthesis phase. The remaining steps described in this section are iteratively executed when the extracted counterplay is looping, and only once otherwise. In the former case, we will refer to the iteration as the *inner-cycle*, to distinguish it from the counterstrategy-guided refinement cycle.

5.1 Candidate Assumptions Computation

Refinements of environment assumptions are computed in four steps: (i) production of two inconsistent Boolean formulae from the counterplay and the unrealizable core, (ii) interpolation between the two Boolean formulae, (iii) translation of the interpolant into LTL, and (iv) negation of the translated interpolant.

Step (i) is executed by the functions **TranslateCounterplayAssumptions** and **TranslateGuarantees** (lines 7–8). The procedure employs the translation scheme in [7] for bounded model checking: it ensures that the obtained Boolean formula is satisfiable if and only if the play taken into account satisfies the LTL formula. The inclusion of assumptions in the counterplay translation is important in yielding an interpolant in the shared alphabet of assumptions and guarantees that explains why the assumptions violate the guarantees. Given a GR(1) formula in φ^θ over \mathcal{V} and a counterplay $\pi_{C,u}$ with state space $S_\pi \in \mathcal{S}$, its translation is a Boolean formula over the domain $\mathcal{V}(S_\pi)$ obtained by replicating every variable $p \in \mathcal{V}$ for every state $s \in S_\pi$; we denote by $p(s)$ the replica of p corresponding to state s , and by $\mathcal{V}(s)$ the subset of $\mathcal{V}(S_\pi)$ containing all the variables referring to state s . This step produces two formulae:

- $[[\pi_{C,u}, \varphi_u^\mathcal{E}]]$, which is a conjunction between the assumptions translation $[[\varphi_u^\mathcal{E}]]$ over $\pi_{C,u}$ and a formula representing the valuation of every $s \in S_\pi$ in $\pi_{C,u}$; the latter is a conjunctive formula containing a literal $p(s)$ (resp. $\neg p(s)$) for every $p \in \mathcal{X} \cup \mathcal{Y}'$ that is true (resp. false) in $\lambda'(s)$ (see end of Sect. 3);
- $[[\varphi_u^\mathcal{S}]]$ which is the guarantees translation over $\pi_{C,u}$.

The translations $[[\varphi_u^\mathcal{E}]]$ and $[[\varphi_u^\mathcal{S}]]$ are given in the extended version of this work [11]. Since by definition a counterplay π_C satisfies the assumptions and violates the guarantees, the formula $[[\pi_{C,u}, \varphi_u^\mathcal{E}]] \wedge [[\varphi_u^\mathcal{S}]]$ is unsatisfiable by construction. Therefore, there exists an interpolant for $[[\pi_{C,u}, \varphi_u^\mathcal{E}]]$ and $[[\varphi_u^\mathcal{S}]]$.

Step (ii) consists of the function **Interpolate** (line 9). The returned interpolant I_u is an over-approximation of $[[\pi_{C,u}, \varphi_u^\mathcal{E}]]$ which by definition implies the negation of $[[\varphi_u^\mathcal{S}]]$: it can be interpreted as a cause of the guarantees not being satisfied by the counterplay, and as such a characterization of a set of counterplays not satisfying the guarantees.

From such interpolant the procedure aims at extracting a set of refinements that fit the GR(1) format. In order to do this, the Boolean to temporal translation requires the interpolant to adhere a specific structure. This is embodied in the notion of *full-separability*. To formally define full-separability, we need first to define state-separability and I/O-separability.

Definition 2 (State-separable interpolant). *An interpolant I_u is said to be state-separable iff it can be expressed as*

$$\bigwedge_{s \in S_u} B_s(\mathcal{V}(s)) \tag{1}$$

where $B_s(\mathcal{V}(s))$ is a Boolean formula either equal to true or expressed over variables in $\mathcal{V}(s)$ only.

We will refer to each $B_s(\mathcal{V}(s))$ as a *state component* of the interpolant. In particular, a state component is equal to *true* if I_u does not use any variables from s . State-separability intuitively means that the subformulae of the interpolant involving a single state are linked by conjunctions. This means that in any model of the interpolant each state component must be itself *true*.

Definition 3 (I/O-separable Boolean expression). *A Boolean expression $B_s(\mathcal{V}(s))$ is said to be I/O-separable if it can be written as a conjunction of two subformulae containing only input and output variables respectively:*

$$B_s(\mathcal{V}(s)) = B_{s,\mathcal{X}}(\mathcal{X}(s)) \wedge B_{s,\mathcal{Y}}(\mathcal{Y}(s)) \tag{2}$$

We call $B_{s,\mathcal{X}}(\mathcal{X}(s))$ and $B_{s,\mathcal{Y}}(\mathcal{Y}(s))$ the *projections* of $B_s(\mathcal{V}(s))$ onto \mathcal{X} and \mathcal{Y} respectively. Any model of an I/O-separable Boolean expression satisfies the projections separately. We can now define full-separability of an interpolant.

Definition 4 (Fully-separable interpolant). *An interpolant is called fully-separable if it is state-separable and each of its state components is I/O-separable.*

An example of a fully-separable interpolant over $\mathcal{X} = \{a, b\}, \mathcal{Y} = \{c, d\}$ and states $S = \{s_0, s_1\}$ is $(a(s_0) \vee b(s_0)) \wedge c(s_0) \wedge \neg b(s_1)$; a non-fully-separable interpolant, instead, is $a(s_0) \vee a(s_1)$, since literals referring to different states are linked via a disjunction.

Remark 1. A particular class of fully-separable interpolants is that of fully conjunctive interpolants, where no disjunctions appear. Whether or not the resulting interpolant is conjunctive depends on the order in which the interpolation algorithm [30] chooses the root clauses for building the unsatisfiability proof. A sufficient condition for obtaining a fully-conjunctive interpolant is that such root clauses be single literals from $[[\pi_{C,u}, \varphi_u^E]]$, and that the pivot variable in each resolution step belong to the shared alphabet of $[[\pi_{C,u}, \varphi_u^E]]$ and $[[\varphi_u^S]]$. (see [11, 30] for details on the interpolation algorithm used).

Step (iii) consists of the function **TranslateInterpolant** (line 14). It converts a fully-separable interpolant $I_u = \bigwedge_{s \in S_u} B_s(\mathcal{V}(s))$ into the LTL formula

$$\begin{aligned} \mathcal{T}(I_u) = & B_{\mathcal{X}}^{init}(\mathcal{X}) \wedge \bigwedge_{s \in S_u} \mathbf{F} \left(B_s(\mathcal{V}) \wedge B_{\text{succ}(s), \mathcal{X}}(\mathbf{X}\mathcal{X}) \right) \wedge \\ & \mathbf{FG} \bigvee_{j=1}^{|S^{loop}|} \left(B_j^{loop}(\mathcal{V}) \wedge \bigwedge_{r=1}^u B_{j,r}^{unr}(\mathcal{V}) \right) \end{aligned} \tag{3}$$

where the expression $B_{\mathcal{X}}^{init}(\mathcal{X})$ is a shorthand for $B_{s^{init},\mathcal{X}}(\mathcal{X})$, $B_j^{loop}(\mathcal{V})$ for $B_{s_j^{loop}}(\mathcal{V})$ and $B_{j,r}^{unr}(\mathcal{V})$ for $B_{s_{j,r}^{unr}}(\mathcal{V})$. Formula (3) is formed from the single state components of I_u by replacing the variables in $\mathcal{V}(s)$ with the corresponding variables in \mathcal{V} and by projecting the components onto the input variables where required by the GR(1) template. The translation consists of three units: a subformula describing the initial state, a conjunction of **F** formulae each containing two consecutive state components, and an **FG** formula; this unit consists of a disjunction over all the looping states, where each disjunct j groups the state components of all the replicas of state s_j^{loop} .

Formula (3) is guaranteed to hold in the counterplay π_C . Intuitively, since I_u is fully-separable by construction, $[[\pi_{C,u}, \varphi_u^\mathcal{E}]]$ implies each state component and its projections onto \mathcal{X} and \mathcal{Y}' . A state component $B_s(\mathcal{V}(s))$ corresponds to a formula $B_s(\mathcal{V})$ satisfied by state s of the counterplay. Therefore, since the initial state satisfies $B^{init}(\mathcal{V})$, π_C satisfies $B_{\mathcal{X}}^{init}(\mathcal{X})$; since there are two consecutive states s and $\text{succ}(s)$ that satisfy $B_s(\mathcal{V}(s))$ and $B_s(\mathcal{V}(\text{succ}(s)))$ respectively, π_C satisfies **F** $(B_s(\mathcal{V}) \wedge B_{\text{succ}(s),\mathcal{X}}(\mathbf{X}\mathcal{X}))$. Finally, for the **FG** subformula, it is sufficient to observe that the looping state j satisfies the formula obtained from the state components referring to s_j^{loop} and $s_{j,r}^{unr}$: since the counterplay remains indefinitely in the looping state, there is a suffix of it where such formula is true for at least one j . Based on these considerations, we prove the following soundness property.

Theorem 1. *Let π_C be a counterplay and $\varphi^\mathcal{E}$ a set of assumptions satisfied in π_C , such that their Boolean translation $[[\pi_{C,u}, \varphi_u^\mathcal{E}]]$ implies I_u , and let I_u be a fully-separable interpolant. Then $\pi_C \models \mathcal{T}(I_u)$.*

The proof is in the extended version [11]. In the case a fully-separable interpolant is not generated from which $\mathcal{T}(I_u)$ can be constructed, the algorithm returns *false* as its candidate assumption. Otherwise, the approach proceeds to step (iv) (function **ExtractDisjuncts**, line 15) producing the candidate refinements by negating (3) and extracting the disjuncts in the resulting formula:

$$\neg B_{\mathcal{X}}^{init}(\mathcal{X}) \vee \bigvee_{s \in S_u} \mathbf{G} \neg \left(B_s(\mathcal{V}) \wedge B_{\text{succ}(s),\mathcal{X}}(\mathbf{X}\mathcal{X}) \right) \vee \mathbf{GF} \bigwedge_{j=1}^{|S^{loop}|} \neg \left(B_j^{loop}(\mathcal{V}) \wedge \bigwedge_{r=1}^u B_{j,r}^{unr}(\mathcal{V}) \right) \quad (4)$$

Each disjunct above is a GR(1) candidate assumption which, by Theorem 1, ensures the exclusion of the counterplay π_C from the models of the assumptions.

5.2 Equivalence Checking and Unrolling

The equivalence checking of the produced candidates and the unrolling of the counterplay (lines 17–24) are only executed in case of a looping counterplay.

Thus in each iteration of the inner-cycle, our procedure checks whether the synthesized assumptions are equivalent to the assumptions ψ_{old} computed in the previous iteration. If not, the looping part of the counterplay is unrolled once (**UnrollCounterplay**, line 20) and the steps in Sect. 5.1–5.2 are repeated. If the equivalence condition is met, the synthesis procedure returns the last set of computed candidates as output.

Counterplay unrolling consists in making the first traversals of looping states explicit. It is achieved by augmenting a counterplay with replicates of the looping states. The number of unrollings is referred to as the *unrolling degree* u . Each unrolling yields a new set of states $S^{unr} = \{s_{1,1}^{unr}, \dots, s_{h,1}^{unr}, \dots, s_{1,u}^{unr}, \dots, s_{h,u}^{unr}\}$. An unrolled looping counterplay has the form $s_{1,1}^{init} s_1^{trans} \dots s_k^{trans} s_{1,1}^{unr} \dots s_{h,1}^{unr} \dots s_{1,u}^{unr} \dots s_{h,u}^{unr} (s_1^{loop} \dots s_k^{loop})^\omega$. Unrolling has two possible effects on the computed interpolant: on one hand, it can introduce new state components in the interpolant, which yield new invariant refinements according to (4); on the other hand, the interpolant can express a more specific characterization of looping states, which corresponds to a weaker fairness refinement in (4). These effects are both observed in our evaluation (see Sect. 7).

6 Convergence

Our procedure is guaranteed to terminate after a finite number of recursive calls. We discuss below the case of all computed interpolants being fully-separable. If not, the procedure terminates with a trivial assumption refinement *false*.

Theorem 2. *Given a satisfiable but unrealizable specification $\langle \phi^E, \phi^S \rangle$ Algorithm 1 terminates with a realizable specification $\langle \phi^{E'}, \phi^S \rangle$.*

To prove this, it is sufficient to show that both the recursion in Algorithm 1 and the iteration over unrollings in Algorithm 2 reach the respective termination conditions. In the following arguments, we will refer to the recursion tree of Algorithm 1. Each node is associated with the candidate assumption tested in one specific call of **CounterstrategyGuidedRefinements**. The root corresponds to the initial assumption; every internal node symbolizes an unrealizable assumptions refinement; the children of an internal node correspond to the alternative refinements that rule out the relevant counterstrategy. The leaves represent alternative realizable assumption refinements returned by the algorithm. We will show that this tree has finite depth and breadth.

Let us consider the number of children n_C of an internal node (the subscript C indicates the counterstrategy computed in that internal node). It consists of the maximum number of refinements that are generated from a single counterstrategy. Assuming that the maximum unrolling degree is finite (we will see that later in this section), denoted $u_{C,MAX}$, the maximum number of refinements generated from C can be computed by counting the maximum number of disjuncts in (4). Suppose $|S_{u_{C,MAX}}|$ denotes the number of distinct states in the unrolled counterplay, then $n_C \leq |S_{u_{C,MAX}}| + 2$: we count one initial condition,

one fairness condition and $|S_{u_C, MAX}|$ invariants. Given that every node has a finite number of children, the breadth of each level in the tree is also finite.

We now consider the depth. The algorithm keeps refining a computed assumption until the property becomes realizable (in case the returned refinement is *false*, then the property is realizable, and therefore the algorithm reaches a *true* leaf). Given the soundness property, at each step every refinement excludes the latest computed counterstrategy; since this counterstrategy satisfies all the previously computed refinements by definition, the new refinement cannot be equivalent to any of the previous refinements along the same branch.

For the above reason, the depth d of the recursion tree is limited by the maximum number of existing GR(1) refinements modulo logical equivalence. The maximum number of initial conditions is $d_{init, MAX} = 2^{2^{|\mathcal{X}|}}$, that is the number of all distinct Boolean expressions over the input variables. The maximum number of invariants is $d_{inv, MAX} = 2^{2^{|\mathcal{V}|+2^{|\mathcal{X}|}}}$; this corresponds to the maximum number of distinct B_s that can be present in the expression (4) times the number of distinct $B_{succ(s), X}$. Finally, the maximum number of distinct fairness assumptions is $d_{fair, MAX} = 2^{2^{|\mathcal{V}|}}$. Therefore, the total depth d is bounded by the sum of these three quantities: $d \leq d_{MAX} = d_{init, MAX} + d_{inv, MAX} + d_{fair, MAX}$.

Given the above, we conclude that the recursion tree is finite. This gives us a worst-case upper bound on the depth d of the recursion, which has a doubly exponential growth over $|\mathcal{V}|$ — a general observation of counterstrategy-guided assumptions refinement strategies. It remains to show that the inner-cycle terminates in finite time. As mentioned in Sect. 5.2, each iteration can provide additional or weaker refinements with respect to the previous iteration. The termination condition holds when the current iteration does not yield new refinements with respect to the previous one. This is reached in the worst case after all distinct GR(1) refinements are generated. The computation is the same as the one for d : $u_{C, MAX} = d_{MAX}$.

7 Evaluation

We apply our approach to two benchmarks presented in [3, 9, 23]: a lift controller and ARM’s AMBA-AHB protocol. The requirements analysis tool RATS_Y [8] is used to check unrealizability and compute counterstrategies. The SAT solver MathSAT [10, 14] is used to compute interpolants. We implemented a translation module for GR(1) specifications and randomly extracted plays into a propositional logic format executable by MathSAT. For each case study, we report the maximum depth and breadth of the recursion tree, and an interpretation of some interesting refinements that are computed. Details are available at [1].

Table 1 provides a summary of both case studies. The columns **In** and **Out** contain the number of input and output variables in the specification alphabet respectively; **A** and **G** contain the number of assumptions and guarantees respectively; **MaxPlay** contains the maximum number of states in a counterplay among all the counterplays used in the refinement process; **MaxUnr** reports the maximum unrolling degree reached in any step of the approach before reaching

the termination condition; **TreeDepth** corresponds to the depth of the recursion tree; **MaxAltRef** is the maximum number of alternative refinements computed to rule out any single counterstrategy (it corresponds to the maximum number of children of an internal node in the recursion tree); **#Ref** shows the total number of refinement sets computed that make the property realizable.

Table 1. Summary of refinement results on benchmarks

Specification	In	Out	A	G	MaxPlay	MaxUnr	TreeDepth	MaxAltRef	#Ref
Lift	3	3	7	12	2	2	1	3	3
AMBA02	7	16	10	66	4	2	3	6	17
AMBA04	11	23	16	97	7	1	2	2	8
AMBA08	19	36	28	157	18	1	7	2	80

7.1 Lift Controller

This case study (also used for controller synthesis problems [3,9]) involves the specification of a system comprising a lift controller. The lift moves between three floors. The environment consists of three buttons, whose states can be *pressed* or *unpressed*; the corresponding state is represented by three binary input variables $\{b_1, b_2, b_3\}$. The controller's state consists of three output variables $\{f_1, f_2, f_3\}$ which indicate at which floor the lift is. The assumptions are:

1. $\varphi_{init}^e = \neg b_1 \wedge \neg b_2 \wedge \neg b_3$
2. $\varphi_{1,i}^e = \mathbf{G}(b_i \wedge f_i \rightarrow \mathbf{X}\neg b_i)$
3. $\varphi_{2,i}^e = \mathbf{G}(b_i \wedge \neg f_i \rightarrow \mathbf{X}b_i)$

for $i \in \{1, 2, 3\}$. They state that the buttons are not pressed in the initial state (1); a pressed button transits to a non-pressed state when the lift arrives at the corresponding floor (2); and the button remains in the pressed state until the lift arrives at that floor (3). The guarantees are:

1. $\varphi_{init}^s = f_1 \wedge \neg f_2 \wedge \neg f_3$
2. $\varphi_1^s = \mathbf{G}(\neg(f_1 \wedge f_2) \wedge \neg(f_2 \wedge f_3) \wedge \neg(f_1 \wedge f_3))$
3. $\varphi_{2,1}^s = \mathbf{G}(f_1 \rightarrow (\mathbf{X}f_1 \vee \mathbf{X}f_2))$
4. $\varphi_{2,2}^s = \mathbf{G}(f_2 \rightarrow (\mathbf{X}f_1 \vee \mathbf{X}f_2 \vee \mathbf{X}f_3))$
5. $\varphi_{2,3}^s = \mathbf{G}(f_3 \rightarrow (\mathbf{X}f_2 \vee \mathbf{X}f_3))$
6. $\varphi_3^s = \mathbf{G}(((f_1 \wedge \mathbf{X}f_2) \vee (f_2 \wedge \mathbf{X}f_3) \vee (f_2 \wedge \mathbf{X}f_1) \vee (f_3 \wedge \mathbf{X}f_2)) \rightarrow (b_1 \vee b_2 \vee b_3))$
7. $\varphi_{4,i}^s = \mathbf{GF}(b_i \rightarrow f_i)$
8. $\varphi_{5,i}^s = \mathbf{GF}f_i$

for $i \in \{1, 2, 3\}$. They state that the lift starts from floor 1 (1); it can never be in two floors at the same time (2); it can move only between consecutive states (3–5), and moves only when at least a button is pressed (6); plays in which the environment keeps a button b_i pressed infinitely and the lift never reaches the corresponding f_i are forbidden (7); and that the lift is required to visit all the

floors infinitely often (8). Given this specification, the fairness guarantee can be satisfied if the environment sets one of its b_i to 1 at least once.

The specification is unrealizable, since when the buttons (environment) stay indefinitely unpressed, the lift (controller) cannot move and therefore $\varphi_{5,2}^s$ and $\varphi_{5,3}^s$ are violated. The unrealizable core consists of the whole set of assumptions and the guarantees φ_{init}^s , $\varphi_{2,1}^s$, φ_3^s and $\varphi_{5,2}^s$. From this core, RATSy computes the counterstrategy π_C in Fig. 1, which consists of a unique play. After translating the unrealizable core over the counterplay, the interpolant is $I_0 = \neg b_1(s_0) \wedge \neg b_2(s_0) \wedge \neg b_3(s_0)$, which corresponds to the GR(1) refinement $\neg b_1 \wedge \neg b_2 \wedge \neg b_3$. The first unrolling is performed yielding the interpolant $I_1 = \neg b_1(s_0) \wedge \neg b_2(s_0) \wedge \neg b_3(s_0) \wedge \neg b_1(s_{1,1}^{unr}) \wedge \neg b_2(s_{1,1}^{unr}) \wedge \neg b_3(s_{1,1}^{unr})$. By translating and negating this interpolant, we obtain the alternative refinements

1. $b_1 \vee b_2 \vee b_3$
2. $\mathbf{G}(\neg b_1 \wedge \neg b_2 \wedge \neg b_3 \rightarrow \mathbf{X}(b_1 \vee b_2 \vee b_3))$
3. $\mathbf{GF}(b_1 \vee b_2 \vee b_3)$

Notice that unrolling results in an interpolant containing an additional state component, thus allowing for more alternative refinements (see Sect. 5.2). Moreover, the new state component refers to an unrolled state, from which a new fairness refinement that is not inferable from I_0 is synthesized. The second unrolling produces equivalent refinements, and thereby the inner-cycle terminates.




Fig. 1. Lift counterstrategy produced by RATSy. The labelling λ' is shown in each state. In this case the lift position plays no role in the environment's choice of next state, therefore $\mathcal{Y}' = \emptyset$.

Every candidate refinement computed by our approach is helpful. Moreover, each one solves the unrealizability problem for the original specification. Refinement (1) does this in a trivial way, since it contradicts the initial assumption contained in the specification. Notice that all the computed refinements force at least one of the buttons to be pressed at some point in any play of the environment. This corresponds to the refinement produced by the approach in [3].

7.2 AMBA-AHB Protocol

The *Advanced High-performance Bus (AHB)* is part of the *Advanced Microcontroller Bus Architecture (AMBA)* specification. It is an open-source communication protocol for on-chip devices through a shared bus. Devices are divided into *masters*, which initiate a communication, and *slaves*, which respond to requests. Multiple masters can request the bus simultaneously, but only one at a time can communicate through it. Masters and slaves constitute the environment, while the system is the bus arbiter implementing the protocol. The specification of the

AHB protocol is provided with RATS_Y. It is a GR(1) description of the protocol in [6], and formalized in [9]. We consider specifications for two, four and eight masters (AMBA02, AMBA04, AMBA08 respectively) which are realizable. To evaluate our approach, we remove the assumption $\mathbf{GF}hready$ as done in [3, 27].

In all the variants, our approach was able to produce refinements that were semantically related to the removed assumption. In the AMBA02 case, one of the refinements is the invariant $\mathbf{G}(hready \vee \mathbf{X}hready)$, which forces *hready* to be true at least every two steps. The other refinements in all the AMBA0x variants involve the variable *hmaster*, which indicates the master that currently owns the bus. These refinements force *hmaster* to change infinitely often. This corresponds to having *hready* equal to *true* infinitely often, since *hready* must be true at any ownership switch according to the protocol [9].

The approach was further tested by extracting different counterplays from the same counterstrategy in the AMBA02 case. Every refinement produced within each synthesis call was helpful. We compared our results to those obtainable through [3, 27] when variables not contained in the interpolant are provided as input. The refinements $\mathbf{GF}(\neg hburst1)$ and $\mathbf{GF}hlock0$ (which are possible outputs of [3, 27] if the user chooses the corresponding templates/variables) remove the first counterstrategy; however neither is helpful, since even after their addition the corresponding set of minimally unfulfillable guarantees is still unrealizable.

8 Discussion

Targeting Unrealizable Cores. The evaluation shows that our approach automatically selects variables that need to be constrained in order to reach realizability. In particular, all the intermediate refinements eliminate precisely a cause of unrealizability, consisting of the set of minimally unfulfillable guarantees from which a counterstrategy has been computed. We note that the returned variables in the AMBA02 example (*hready* and *hbusreq1*) are a subset of the variables that the authors in [3] suggest to use in order to instantiate the refinement templates.

Helpfulness of intermediate refinements is a desirable condition for reducing the convergence rate of the algorithm. When this holds, then the expected tree depth d (see Sect. 6) is reduced to $O(n_g)$, where n_g is the number of minimally unfulfillable subsets of guarantees. The application of our approach on the case studies consistently supports the attainment of this condition.

Number of Unrollings. We further define an upper bound to the number of unrollings needed to reach the termination condition as of Sect. 5.2. Every unrolling iteration produces an interpolant which is either the same as the previous iteration, or contains the description of one more state in the counterplay. In the worst case, without unrolling the interpolant describes just the initial state; after the first unrolling it contains a state component for the first transient state s_1^{trans} ; it is iteratively strengthened by one more state component until it describes all the transient states and the first replica of the unrolled states $s_{j,1}^{unr}$ for each j . In the following unrolling step, the interpolant contains

the component of $s_{1,2}^{unr}$, which is such that $\lambda'(s_{1,2}^{unr}) = \lambda'(s_{1,1}^{unr})$: interpolation produces an equivalent refinement to the previous step, and therefore the procedure terminates. The maximum number of unrollings before reaching the termination condition is: $u_{C,MAX} = |S|$ where $|S|$ is the number of states in π_C .

Comparison with Existing Approaches. Our approach extracts weaker refinements than those of [3]. The reason is that [3] uses templates that are true over all paths of a counterstrategy, while our approach requires them to be true in a single counterplay. More specifically, an invariant template used in [3] has the form $\mathbf{G}\neg q \vee \mathbf{X} \left(\bigwedge_{q' \in Next(q)} \neg q' \right)$, where q and q' indicate states in a counterstrategy and $Next(q)$ is the set of successor states of q ; our approach extracts invariants of the form $\mathbf{G}\neg q \vee \neg q'$ for a $q' \in Next(q)$, which is implied by the former one, provided that they use the same variables set for q and q' .

We notice that in principle our approach may generate assumptions containing only output variables. This happens if some state component in the interpolant contains only output variables. Those are valid GR(1) formulae according to the definition, although hardly interpretable as constraints on the environment. Existing approaches circumvent the problem by allowing only input variables in their refinements [3, 4]: however, in this way valid assumptions are also excluded. In our AMBA04 case study, one of the computed assumptions was $\mathbf{G}((\neg hmaster0 \wedge hbusreq1) \rightarrow \mathbf{X}(\neg hbusreq1))$, where $hmaster0 \in \mathcal{Y}$ and $hbusreq1 \in \mathcal{X}$. This assumption would not have been computed with that restriction.

9 Conclusions

We presented an interpolation-based approach for synthesizing weak environment assumptions for GR(1) specifications. Our approach exploits the information in counterstrategies and unrealizable cores to compute assumptions that directly target the cause of unrealizability. Compared to closely related approaches [3, 27], our algorithm does not require the user to provide the set of variables upon which the assumptions are constructed. The case study applications show that our approach implicitly performs a variable selection that targets an unrealizable core, allowing for a quicker convergence to a realizable specification.

The final set of refinements is influenced by the choice of counterplay. We are investigating in our current work the effect of and criteria over the counterplay selection particularly on the full-separability of interpolants. Furthermore, since interpolants are over-approximations of the counterplays, the final specification is an under-approximation. In future work, we will explore the use of witnesses (winning strategies for the system) to counteract this effect. Finally, the applicability of our approach depends on the separability properties of the computed interpolants: further investigation is needed to characterize the conditions under which an interpolation algorithm returns fully-separable interpolants.

Acknowledgments. The support of the EPSRC HiPEDS CDT (EP/L016796/1) and Imperial College Junior Research Fellowship is gratefully acknowledged. We also thank our reviewers for their insightful comments and suggestions.

References

1. <https://github.com/davidecavezza/TACAS17CaseStudies>
2. Alrajeh, D., Kramer, J., van Lamsweerde, A., Russo, A., Uchitel, S.: Generating obstacle conditions for requirements completeness. In: ACM/IEEE 34th International Conference on Software Engineering, pp. 705–715 (2012)
3. Alur, R., Moarref, S., Topcu, U.: Counter-strategy guided refinement of GR(1) temporal logic specifications. In: Formal Methods in Computer-Aided Design, pp. 26–33 (2013)
4. Alur, R., Moarref, S., Topcu, U.: Pattern-based refinement of assume-guarantee specifications in reactive synthesis. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 501–516. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-46681-0_49](https://doi.org/10.1007/978-3-662-46681-0_49)
5. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)
6. ARM Ltd.: AMBA-AHB specification. <https://developer.arm.com/products/architecture/amba-protocol>
7. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. *Adv. Comput.* **58**, 117–148 (2003)
8. Bloem, R., Cimatti, A., Greimel, K., Hofferek, G., Könighofer, R., Roveri, M., Schuppan, V., Seeber, R.: RATSYS – a new requirements analysis tool with synthesis. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 425–429. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-14295-6_37](https://doi.org/10.1007/978-3-642-14295-6_37)
9. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Sa’Ar, Y.: Synthesis of reactive(1) designs. *J. Comput. Syst. Sci.* **78**(3), 911–938 (2012)
10. Bruttomesso, R., Cimatti, A., Franzén, A., Griggio, A., Sebastiani, R.: The MATHSAT 4 SMT solver. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 299–303. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70545-1_28](https://doi.org/10.1007/978-3-540-70545-1_28)
11. Cavezza, D.G., Alrajeh, D.: Interpolation-based GR(1) assumptions refinement. CoRR abs/1611.07803 (2016), <http://arXiv.org/abs/1611.07803>
12. Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Environment assumptions for synthesis. In: Breugel, F., Chechik, M. (eds.) CONCUR 2008. LNCS, vol. 5201, pp. 147–161. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-85361-9_14](https://doi.org/10.1007/978-3-540-85361-9_14)
13. Cimatti, A., Roveri, M., Schuppan, V., Tchaltsev, A.: Diagnostic information for realizability. In: Logozzo, F., Peled, D.A., Zuck, L.D. (eds.) VMCAI 2008. LNCS, vol. 4905, pp. 52–67. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78163-9_9](https://doi.org/10.1007/978-3-540-78163-9_9)
14. Cimatti, A., Griggio, A., Sebastiani, R.: Efficient interpolant generation in satisfiability modulo theories. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 397–412. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78800-3_30](https://doi.org/10.1007/978-3-540-78800-3_30)
15. Cobleigh, J.M., Giannakopoulou, D., Păsăreanu, C.S.: Learning assumptions for compositional verification. In: Garavel, H., Hatcliff, J. (eds.) TACAS 2003. LNCS, vol. 2619, pp. 331–346. Springer, Heidelberg (2003). doi:[10.1007/3-540-36577-X_24](https://doi.org/10.1007/3-540-36577-X_24)
16. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *J. Symb. Logic* **22**(03), 269–285 (1957)

17. Degiovanni, R., Alrajeh, D., Aguirre, N., Uchitel, S.: Automated goal operationalisation based on interpolation and SAT solving. In: ACM/IEEE 36th International Conference on Software Engineering, pp. 129–139 (2014)
18. D’Silva, V., Kroening, D., Weissenbacher, G.: A survey of automated techniques for formal software verification. *IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst.* **27**(7), 1165–1178 (2008)
19. Esparza, J., Kiefer, S., Schwoon, S.: Abstraction refinement with Craig interpolation and symbolic pushdown systems. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 489–503. Springer, Heidelberg (2006). doi:[10.1007/11691372_35](https://doi.org/10.1007/11691372_35)
20. Gheerbrant, A., Cate, B.: Craig interpolation for linear temporal languages. In: Grädel, E., Kahle, R. (eds.) CSL 2009. LNCS, vol. 5771, pp. 287–301. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-04027-6_22](https://doi.org/10.1007/978-3-642-04027-6_22)
21. Gheorghiu Bobaru, M., Păsăreanu, C.S., Giannakopoulou, D.: Automated assume-guarantee reasoning by abstraction refinement. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 135–148. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70545-1_14](https://doi.org/10.1007/978-3-540-70545-1_14)
22. Kamide, N.: Interpolation theorems for some variants of LTL. *Rep. Math. Logic* **50**, 3–30 (2015)
23. Könighofer, R., Hofferek, G., Bloem, R.: Debugging formal specifications using simple counterstrategies. In: Formal Methods in Computer-Aided Design, pp. 152–159 (2009)
24. Krajíček, J.: Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *J. Symb. Logic* **62**(02), 457–486 (1997)
25. van Lamsweerde, A.: Requirements Engineering - From System Goals to UML Models to Software Specifications. Wiley, Hoboken (2009)
26. van Lamsweerde, A., Letier, E.: Handling obstacles in goal-oriented requirements engineering. *IEEE Trans. Softw. Eng.* **26**(10), 978–1005 (2000)
27. Li, W., Dworkin, L., Seshia, S.A.: Mining assumptions for synthesis. In: ACM/IEEE 9th International Conference on Formal Methods and Models for Codesign, pp. 43–50 (2011)
28. Li, W., Sadigh, D., Sastry, S.S., Seshia, S.A.: Synthesis for human-in-the-loop control systems. In: 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, pp. 470–484 (2014)
29. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems. Springer, Heidelberg (1992)
30. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 1–13. Springer, Heidelberg (2003). doi:[10.1007/978-3-540-45069-6_1](https://doi.org/10.1007/978-3-540-45069-6_1)
31. Păsăreanu, C., Dwyer, M., Huth, M.: Assume-guarantee model checking of software: a comparative case study. In: Dams, D., Gerth, R., Leue, S., Massink, M. (eds.) SPIN 1999. LNCS, vol. 1680, pp. 168–183. Springer, Heidelberg (1999). doi:[10.1007/3-540-48234-2_14](https://doi.org/10.1007/3-540-48234-2_14)
32. Sohail, S., Somenzi, F.: Safety first: a two-stage algorithm for the synthesis of reactive systems. *Int. J. Softw. Tools for Technol. Transfer* **15**(5–6), 433–454 (2013)
33. Sohail, S., Somenzi, F., Ravi, K.: A hybrid algorithm for LTL games. In: Logozzo, F., Peled, D.A., Zuck, L.D. (eds.) VMCAI 2008. LNCS, vol. 4905, pp. 309–323. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78163-9_26](https://doi.org/10.1007/978-3-540-78163-9_26)