

# Fairness, Assumptions, and Guarantees for Extended Bounded Response LTL+P synthesis

Alessandro Cimatti<sup>1</sup>[0000-0002-1315-6990], Luca Geatti<sup>1,2</sup>[0000-0002-7125-787X],  
Nicola Gigante<sup>3</sup>[0000-0002-2254-4821], Angelo Montanari<sup>2</sup>[0000-0002-4322-769X],  
and Stefano Tonetta<sup>1</sup>[0000-0001-9091-7899]

<sup>1</sup> Fondazione Bruno Kessler, Trento, Italy

{cimatti,lgeatti,tonettas}@fbk.eu

<sup>2</sup> University of Udine, Italy

{luca.geatti,angelo.montanari}@uniud.it

<sup>3</sup> Free University of Bozen-Bolzano, Italy

nicola.gigante@unibz.it

**Abstract.** Realizability and reactive synthesis from temporal logics are fundamental problems in the formal verification field. The complexity of the Linear-time Temporal Logic with Past (LTL+P) case led to the definition of fragments with lower complexities and simpler algorithms. Recently, the logic of Extended Bounded Response LTL+P ( $LTL_{EBR}+P$ ) has been introduced. It allows one to express any safety language definable in LTL and it is provided with an efficient, fully-symbolic algorithm for reactive synthesis.

In this paper, we extend  $LTL_{EBR}+P$  with fairness conditions, assumptions, and guarantees. The resulting logic, called GR-EBR, preserves the main strength of  $LTL_{EBR}+P$ , that is, efficient realizability, and makes it possible to specify properties beyond safety. We study the problem of reactive synthesis for GR-EBR and devise a fully-symbolic algorithm that reduces it to a number of safety subproblems. To ensure soundness and completeness, we propose a general framework for safety reductions in the context of realizability of (fragments of) LTL+P. The experimental evaluation shows the feasibility of the approach.

**Keywords:** Realizability · Temporal Logic · Symbolic Automata

## 1 Introduction

One of the most important problems in formal methods and requirement analysis is establishing whether a specification over a set of controllable and uncontrollable actions is *implementable* (or *realizable*), that is, whether there exists a controller that chooses the value of the controllable actions and satisfies the specification, no matter what the values of uncontrollable actions are. This problem has been formalized in the literature under the name of *realizability* [3]. The very close problem of *reactive synthesis* aims at synthesizing such a controller, whenever the specification is realizable. Usually, these problems are modelled as

two-player games between Environment, who tries to violate the specification, and Controller, who tries to fulfill it. Realizability is known to have a very high worst-case complexity. In particular, it has a non-elementary lower bound for S1S specifications [2], and it is 2EXPTIME-complete for LTL specifications [22, 23].

In order to apply realizability and reactive synthesis in real-world scenarios, research has focused on the identification of fragments of logics like S1S and LTL, with a limited expressive power, for which realizability can be solved efficiently.

A well-known example is *Generalized Reactivity(1)* logic (GR(1), for short) [1]. In this fragment, a specification is syntactically partitioned into *assumptions* about the environment and *guarantees* for the controller. Both of them are either Boolean formulas ( $\alpha$ ) or safety formulas ( $G\alpha$ ) or conjunctions of recurrence formulas ( $\bigwedge_{i=1}^n GF\alpha_i$ ). The dichotomy between *assumptions* and *guarantees* reflects the way a system engineer usually formalizes system’s requirements, which is summarized by the following sentence: “*the controller has to behave in conformance to the guarantees, under the given assumptions on the environment*”.

On a different direction, other approaches focused on safety fragments of LTL [4, 25]. In particular, *Extended Bounded Response LTL* ( $LTL_{EBR}$ , for short) is a safety fragment of LTL+P that allows for a fully symbolic compilation of formulas into deterministic automata. Such a feature contributes to a great improvement in solving time for the synthesis problem. Moreover,  $LTL_{EBR}$  has a well-established expressiveness:  $LTL_{EBR}$  can define exactly the set of safety languages definable in LTL.

*Contributions* The main contributions of this paper are the following ones. First, we introduce GR-EBR, an extension of  $LTL_{EBR}$  that admits: (i) fairness conditions, in particular, conjunctions of *recurrence formulas*, that is,  $\bigwedge_i GF\alpha_i$ , forcing each formula  $\alpha_i$  to be true infinitely often; (ii) assumptions/guarantees in the form of an  $LTL_{EBR}$  formula augmented with fairness conditions. In addition to be able to express any  $LTL_{EBR}$  formula, and, consequently, any safety property definable in LTL, GR-EBR allows also for the definition of properties beyond the safety fragment, like, for instance,  $G(p) \rightarrow G(q)$ .

Second, we devise a novel framework for deriving *complete* safety reductions in the context of realizability of (fragments of) LTL. A notable feature of the framework is that it provides a link to safety reductions for the model checking problem and proves that if a reduction is complete for model checking, then it is also complete for realizability. On one hand, this allows one to reason on Kripke structures instead of on strategies, which is simpler; on the other hand, it enables the use of some reductions already exploited in model checking for realizability, provided that they conform to the framework.

Third, the proposed framework is used to derive a *complete* safety reduction for the realizability problem of GR-EBR. A crucial property of the algorithm is that the realizability check of each safety sub-problem is performed in a *fully symbolic* way, thus retaining the distinctive feature of  $LTL_{EBR}$ .

Last but not least, we provide an implementation of the algorithm as a prototype tool called GRACE (*GR-ebr reAlizability ChEcker*). The experimental evaluation shows good performance against tools for full LTL+P synthesis.

*Related work* GR(1) has been introduced in [1, 21]. It is known that GR(1) is a good candidate for writing specifications of real-world scenarios, with a relatively low complexity: the realizability problem can be solved with at most a quadratic number of symbolic steps in the size of the specification [1]. On the other hand, GR(1) presents some restrictions that limit its use as a specification language: (i) safety assumptions/guarantees are either Boolean formulas or formulas of the form  $G\alpha$ , where the only temporal operator admitted in  $\alpha$  is the *next* operator  $X$ ; (ii) assumptions are syntactically constrained to be formulas *controlled* by Environment, in the sense that the variables inside the *next* operators of the safety part of the assumptions must be *uncontrollable*. In GR-EBR we relax that syntactical restrictions of GR(1): for example, the safety assumptions and guarantees can be any arbitrary  $LTL_{EBR}$  formula, like, for instance,  $G(r \rightarrow F^{[0,10]}g)$ . For this reason, GR-EBR can be considered an extension not only of  $LTL_{EBR}$ , but also of GR(1). On the semantic side, the comparison is more problematic. On the one hand, all (standard) realizability problems for GR(1) specifications are definable in LTL+P [1] and also in GR-EBR. On the other hand, we do not know whether GR-EBR is able to express more properties than GR(1). Our conjecture is that this is the case. Take for instance the bounded-response property  $G(r \rightarrow F^{[0,k]}g)$ : it is easily expressible in GR-EBR, but we see no way it could be definable in GR(1) without introducing additional variables (that would maintain realizability but not language equivalence).

*Bounded synthesis* [9, 13] belongs to the class of *Safrless techniques* [17], and it consists in bounding the number of times Controller is forced to visit a rejecting state of a Universal co-Büchi automaton (UCW, for short) for the initial formula. This corresponds to a safety automaton, which can be either (i) made deterministic by a suitable generalization of the classical subset construction [7, 11], or (ii) encoded into a constraint system [9, 13] (*e.g.*, SAT- or SMT-based) which bounds also the *size* of a candidate controller (this also allows one to tackle undecidable problems, for instance in the case of distributed or parametric synthesis). Both choices work for the whole class of UCW, and thus for full LTL. A significant drawback of such an approach is that the UCW, which can be exponentially larger than the initial specification, is explicitly represented. Moreover, in the first case, the algorithm for the determinization turns out to be quite complex, since each state of the resulting automaton is actually a *function*. This can also result into a very large state space, that can be tackled by exploiting either *antichains* [11] or BDDs [7]. In contrast, as we will see, we define a reduction tailored to GR-EBR formulas that allows us to exploit the  $LTL_{EBR}$  transformations introduced in [4] for a *fully symbolic* mapping of the initial formula directly into a sequence of symbolic safety automata. In particular, we never build any explicit-state automaton and we avoid the subsequent use of determinization algorithms.

*Organization* The rest of the paper is organized as follows. In Sec. 2, we introduce the notation and provide the basic definitions. In Sec. 3, we define the logic GR-EBR and give an example of GR-EBR specification. The framework for deriving complete reductions is presented in Sec. 4. In Sec. 5 we describe the

algorithm for the realizability of GR-EBR specifications. The outcomes of the experimental evaluation are reported in Sec. 6. Finally, in Sec. 7, we point out some interesting future research directions.

## 2 Preliminaries

### 2.1 Temporal Logics

Linear Temporal Logic with Past (LTL+P) is a modal logic interpreted over infinite state sequences. Let  $\Sigma$  be a set of propositions. LTL+P formulas are inductively defined as follows:

$$\phi := p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid X\phi \mid \phi_1 U \phi_2 \mid Y\phi \mid \phi_1 S \phi_2$$

where  $p \in \Sigma$ . Temporal operators can be subdivided into the *future operators*, *next* (X) and *until* (U), and *past operators*, *yesterday* (Y) and *since* (S). We define the following common abbreviations (where  $\top$  stands for any tautology such as  $p \vee \neg p$ ): (i) *release*:  $\phi_1 R \phi_2 \equiv \neg(\neg\phi_1 U \neg\phi_2)$ ; (ii) *eventually*:  $F\phi_1 \equiv \top U \phi_1$ ; (iii) *globally*:  $G\phi_1 \equiv \neg F\neg\phi_1$ ; (iv) *once*:  $O\phi_1 \equiv \top S \phi_1$ ; (v) *historically*:  $H\phi_1 \equiv \neg O\neg\phi_1$ . LTL+P formulas are interpreted over infinite state sequences (or  $\omega$ -words)  $\pi \in (2^\Sigma)^\omega$ . We call *language* a set of  $\omega$ -words. We write  $\pi \models \phi$  to denote the fact that the state sequence  $\pi$  is a *model* (or *satisfies*) the formula  $\phi$ . We refer to [4] for the semantics of the LTL+P operators. With  $|\phi|$ , we refer to the *size* of the formula  $\phi$ , defined as the number of symbols in it. We define the *language* of  $\phi$ , written  $\mathcal{L}(\phi)$ , as the set of all and only the models of  $\phi$ .

An important class of languages is the class of safety properties, that express the fact that “something bad never happens”. Formally, we define a *safety property* (or *safety language*) as a language for which it holds that, for any  $\omega$ -word that does not belong to language, there exists a finite prefix of it such that all its continuations do not belong to the language as well. A formula  $\phi$  is called a *safety formula* if  $\mathcal{L}(\phi)$  is a safety language. Recently, Cimatti *et al.* [4] introduced a subset of LTL+P, called *Extended Bounded Response* LTL+P, which expresses exactly the safety properties that can be defined in LTL+P [5], and gave a symbolic procedure to turn formulas of this fragment into symbolic automata.

**Definition 1 (The logic  $LTL_{EBR+P}$  [4]).** *Let  $a, b \in \mathbb{N}$ . An  $LTL_{EBR+P}$  formula  $\chi$  is inductively defined as follows:*

$$\begin{array}{ll} \eta := p \mid \neg\eta \mid \eta_1 \vee \eta_2 \mid Y\eta \mid \eta_1 S \eta_2 & \text{Pure Past Layer} \\ \psi := \eta \mid \neg\psi \mid \psi_1 \vee \psi_2 \mid X\psi \mid \psi_1 U^{[a,b]} \psi_2 & \text{Bounded Future Layer} \\ \phi := \psi \mid \phi_1 \wedge \phi_2 \mid X\phi \mid G\phi \mid \psi R \phi & \text{Future Layer} \\ \chi := \phi \mid \chi_1 \vee \chi_2 \mid \chi_1 \wedge \chi_2 & \text{Boolean Layer} \end{array}$$

We define the *bounded until* operator  $\psi_1 U^{[a,b]} \psi_2$  as a *shortcut* of the formula  $\bigvee_{i=a}^b (X^i \psi_2 \wedge \bigwedge_{j=0}^{i-1} X^j \psi_1)$ , where  $X^i \phi := X_{(1)} \dots X_{(i)} \phi$ . We define  $LTL_P$  (the *pure past fragment of LTL+P*) as the set of all the formulas belonging to the *Pure*

*Past Layer* of Def. 1, respectively. With some abuse of notation, we will denote with the symbol of the logic (e.g., LTL+P or LTL<sub>EBR</sub>+P) also the *set* of all the formulas of the respective logic.

## 2.2 Automata

Temporal logic has a strong relation with automata on infinite words [24]. Since in the following we will work only with symbolic representations, we give here the definition of symbolic automata. It is well-known that the symbolic representation can be exponentially more succinct than the explicit-state one.

**Definition 2 (Symbolic Automaton on Infinite Words).** A symbolic automaton on infinite words over the alphabet  $\Sigma$  is a tuple  $\mathcal{A} = (V, I, T, \alpha)$ , such that (i)  $V = X \cup \Sigma$ , where  $X$  is a set of state variables and  $\Sigma$  is a set of input variables, (ii)  $I(X)$  and  $T(X, \Sigma, X')$ , with  $X' = \{x' \mid x \in X\}$ , are Boolean formulas which define the set of initial states and the transition relation, respectively, and (iii)  $\alpha(X)$  is an LTL+P formula over the variables in  $X$  which defines the accepting condition.

**Definition 3 (Languages of Symbolic Automata).** An  $\omega$ -word (or simply a word)  $\sigma = \langle \sigma_0, \sigma_1, \dots \rangle$  is an infinite sequence of letters in  $\Sigma$ . A run  $\tau = \langle \tau_0, \tau_1, \dots \rangle$  is an infinite sequence of states (i.e., evaluations of the variables in  $X$ ) that are in relation with respect to  $T$  (i.e., such that any two consecutive evaluations satisfy the formula  $T$ ). A run  $\tau$  is induced by the word  $\sigma$  iff  $\tau_0 \models I$  and  $(\tau_i, \sigma_i, \tau_{i+1}) \models T$ , for all  $i \geq 0$ . We say that  $\mathcal{A}$  is deterministic iff there exists exactly one trace induced by  $\sigma$ , for each  $\sigma \in \Sigma^\omega$ . A word  $\sigma$  is accepted by  $\mathcal{A}$  iff there exists an accepting run induced by  $\sigma$  in  $\mathcal{A}$ . The language of  $\mathcal{A}$ , denoted with  $\mathcal{L}(\mathcal{A})$ , is the set of all and only the words accepted by  $\mathcal{A}$ .

We will refer to three important classes of accepting conditions: (i) *safety*:  $\alpha(X) := \mathbf{G}\beta$ ; (ii) *Reactivity(1)*:  $\alpha(X) := \mathbf{GF}\beta \rightarrow \mathbf{GF}\beta'$ ; (iii) *Generalized Reactivity(1)*:  $\alpha(X) := \bigwedge_{i=1}^m \mathbf{GF}\beta_i \rightarrow \bigwedge_{j=1}^n \mathbf{GF}\beta'_j$ ; where each  $\beta, \beta', \beta_i, \beta'_j \in \text{LTL}_P$ .

## 2.3 Model Checking, Realizability, and Synthesis

A Kripke structure is a tuple  $M = (\Sigma, Q, I, T, L)$  where: (i)  $\Sigma$  is the input alphabet, (ii)  $Q$  is the (finite) set of states, (iii)  $I \subseteq Q$  is the set of initial states, (iv)  $T \subseteq Q \times Q$  is a *complete* transition relation, and (v)  $L : Q \rightarrow 2^\Sigma$  is the labeling function that assigns to each state the set of atoms in  $\Sigma$  that are true in that state. We denote with  $|M|$  the number of states in  $M$ , i.e.,  $|Q|$ . Given a path  $\pi := \langle q_0, q_1, \dots \rangle$  in  $M$ , we denote with  $L(\pi)$  the sequence  $\langle L(q_0), L(q_1), \dots \rangle$ . The path  $\pi$  is called *initialized* iff  $q_0 \in I$ . The model checking problem takes as input a Kripke structure and a temporal formula, and asks to find whether all the initialized traces of the former satisfy the latter.

**Definition 4 (The model checking problem).** Given a Kripke structure  $M$  and a linear temporal formula  $\phi$ , the model checking problem is the problem of finding whether all the initialized traces  $\pi$  of  $M$  are such that  $L(\pi) \models \phi$ , written  $M \models A\phi$  (where  $A$  is the “for all paths” operator of CTL).

Realizability and reactive synthesis are in some sense more ambitious problems than model checking, since they aim to find whether a given temporal formula  $\phi$  over two sets  $\mathcal{U}$  and  $\mathcal{C}$  of uncontrollable and controllable variables, respectively, is implementable and, if this is the case, to synthesize a possible *implementation*. Usually, realizability is modeled as a two-player game between Environment, who tries to violate the specification and Controller, who tries to fulfill it. In this setting, an implementation of the specification is represented by a *strategy*.

**Definition 5 (Strategies and Languages of Strategies).** *Let  $\mathcal{U}$  and  $\mathcal{C}$  be two disjoint sets of input (or uncontrollable) and output (or controllable) variables, respectively. A strategy  $g$  is a function  $g : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$ . We define the language of the strategy  $g$ , denoted as  $\mathcal{L}(g)$ , as the set of all and only the sequences  $\langle (U_0 \cup C_0), (U_1 \cup C_1), \dots \rangle$  such that  $U_i \in 2^{\mathcal{U}}$  and  $C_i = g(\langle U_0, \dots, U_i \rangle)$ , for all  $i \geq 0$ .*

**Definition 6 (Realizability and Synthesis for LTL).** *Let  $\phi$  be a temporal formula over the alphabet  $\Sigma = \mathcal{U} \cup \mathcal{C}$ , where  $\mathcal{U}$  is the set of input variables,  $\mathcal{C}$  the set of output variables, and  $\mathcal{U} \cap \mathcal{C} = \emptyset$ . We say that  $\phi$  is realizable if and only if there exists a strategy  $g : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$  such that  $\mathcal{L}(g) \subseteq \mathcal{L}(\phi)$ . If  $\phi$  is realizable, the synthesis problem is the problem of computing such a strategy.*

The strategies which we are mainly interested in are the ones that can be represented finitely. In the literature, there are two main (and equivalent) representations for finite strategies, that is, *Mealy machines* and *Moore machines*. In this paper, we are mainly interested in the first ones.

**Definition 7 (Mealy Machine).** *A Mealy machine is a tuple  $M = (\Sigma_{\mathcal{U}}, \Sigma_{\mathcal{C}}, Q, q_0, \delta)$  such that: (i)  $\Sigma_{\mathcal{U}}$  and  $\Sigma_{\mathcal{C}}$  are the input and output alphabets, respectively; (ii)  $Q$  is the (finite) set of states and  $q_0$  is the initial state; (iii)  $\delta : Q \times \Sigma_{\mathcal{U}} \rightarrow \Sigma_{\mathcal{C}} \times Q$  is the total transition function. We say that an infinite word  $\sigma = \langle \sigma_0, \sigma_1, \dots \rangle \in (\Sigma_{\mathcal{U}} \cup \Sigma_{\mathcal{C}})^\omega$  is accepted by  $M$  iff there exists a trace  $\langle (q_0, \sigma_0), (q_1, \sigma_1), \dots \rangle \in (Q \times (\Sigma_{\mathcal{U}} \cup \Sigma_{\mathcal{C}}))^\omega$  such that  $\delta(q_i, \sigma_i \cap \Sigma_{\mathcal{U}}) = (\sigma_i \cap \Sigma_{\mathcal{C}}, q^{i+1})$ , for all  $i \geq 0$ . We define the language of  $M$ , written as  $\mathcal{L}(M)$ , as the set of all the infinite words accepted by  $M$ .*

A fundamental feature is the *small model property* for realizability of LTL+P [11, 17, 22], which ensures that each realizable LTL+P formula has at least a finitely representable strategy.

**Proposition 1 (Small model property of LTL+P [22]).** *Let  $\phi$  be an LTL+P formula and  $n = |\phi|$ . If  $\phi$  is realizable by a strategy  $g$ , then there exists a Mealy machine  $M_g$  such that (i)  $M_g$  has at most  $2^{2^{c \cdot n}}$  states, for some constant  $c \in \mathbb{N}$ , and (ii)  $\mathcal{L}(M_g) \subseteq \mathcal{L}(\phi)$ .*

### 3 LTL<sub>EBR</sub>+P with fairness, assumptions, and guarantees

In this section, we extend LTL<sub>EBR</sub>+P (see Def. 1) with fairness conditions (*i.e.*, of type GF $\alpha$ ), assumptions and guarantees (that correspond to the antecedent

and the consequent of a logical implication). The syntax of the resulting logic, called GR-EBR, is the following.

**Definition 8 (The logic GR-EBR).** *The GR-EBR logic comprises all and only those formulas that can be written in the following form:*

$$(\psi_{ebr}^1 \wedge \bigwedge_{i=1}^m \text{GF}\alpha_i) \rightarrow (\psi_{ebr}^2 \wedge \bigwedge_{j=1}^n \text{GF}\beta_j)$$

where  $m, n \in \mathbb{N}$ ,  $\psi_{ebr}^1, \psi_{ebr}^2 \in \text{LTL}_{\text{EBR}+\text{P}}$  and  $\alpha_i, \beta_j \in \text{LTL}_{\text{P}}$ , for each  $i, j \in \mathbb{N}$ .

### 3.1 Expressiveness of GR-EBR

Each  $\text{LTL}_{\text{EBR}+\text{P}}$  formula  $\phi$  is a GR-EBR formula as well. In fact,  $\phi \equiv (\top \wedge \top) \rightarrow (\phi \wedge \top) \in \text{GR-EBR}$ . It follows that any safety language definable in  $\text{LTL}+\text{P}$  is definable in GR-EBR as well. In addition, GR-EBR is *strictly* more expressive than  $\text{LTL}_{\text{EBR}+\text{P}}$ , since the former can express also non-safety properties, like  $\text{G}(p) \rightarrow \text{G}(q)$ .

Consider the temporal hierarchy defined by Manna and Pnueli in [19]. The Reactivity class is defined as the set of all and only those languages definable by formulas of type  $\bigwedge_i (\text{GF}\alpha_i \rightarrow \text{GF}\beta_i)$  where each  $\alpha_i$  and each  $\beta_i$  are pure-past  $\text{LTL}+\text{P}$  formulas. It is known that  $\text{LTL}+\text{P}$  is expressively equivalent to the Reactivity class. Moreover, if we fix the number of conjuncts of the formula above to be  $\mathbb{N}$ , that the resulting class (called Reactivity( $\mathbb{N}$ )) strictly contains Reactivity( $\mathbb{N}-1$ ) and is strictly contained in Reactivity( $\mathbb{N}+1$ ). Compared to this classification, we have that GR-EBR is at least as expressive as the Reactivity(1) class, since each formula of type  $\text{GF}(\alpha) \rightarrow \text{GF}(\beta)$  belongs to GR-EBR. However, the exact expressiveness of GR-EBR is still unknown.

On a more practical side, we found that some benchmarks of SYNTCOMP [15], like *simple\_arbiter\_N* (for each  $N \in \{2, 4, 6, 8, 10, 12\}$ ) and also *escalator\_bidirectional*, can be translated in GR-EBR with minor rewritings.

### 3.2 An Example

We take the example proposed in [4] and we extend it with fairness, assumptions and guarantees. Suppose that we want to synthesize an arbiter that, given a request from client  $i$  (for some  $i \in \{1, \dots, n\}$ ) in the environment, assigns the grant to the corresponding client, in such a way to guarantee the following properties: (1) *bounded response*: the grant is assigned at most  $k$  time units, for some  $k > n$ , after the request is issued; (2) *mutual exclusion*: the arbiter can assign a grant at most to one client at a time. The conjunction of the previous two requirements form the guarantees for the controller. The assumptions for the environment are the following: (1) initially, there are no requests; (2) if a request is issued at time  $i$ , then it cannot be issued until time  $i + k$ ; (3) there are infinitely many requests from each client.

In order to write a specification of the arbiter, we can model the requests for the  $n$  clients with the (uncontrollable) variables  $r_1, \dots, r_n$ . Similarly, the grant corresponding to the request  $r_i$  can be modeled with the (controllable) variable  $g_i$ , for each  $i \in \{1, \dots, n\}$ . The assumption for the environment corresponds to the  $\text{LTL}_{\text{EBR}+\text{P}}$  formula  $\phi_e$  defined as follows:

$$\bigwedge_{i=1}^n \neg r_i \wedge \bigwedge_{i=1}^n \text{G}(r_i \rightarrow \text{G}^{[1,k]}\neg r_i) \wedge \bigwedge_{i=1}^n \text{GF}r_i$$

The guarantees for the controller correspond to the  $\text{LTL}_{\text{EBR}+\text{P}}$  formula  $\phi_c$  defined as follows:

$$\bigwedge_{i=1}^n \text{G}(r_i \rightarrow \text{F}^{[0,k]}g_i) \wedge \text{G}\left(\bigwedge_{1 \leq i < j \leq n} \neg(g_i \wedge g_j)\right)$$

The overall specification is  $\phi_e \rightarrow \phi_c$  and syntactically belongs to  $\text{GR-EBR}$ .

Our goal is to solve the realizability problem for  $\text{GR-EBR}$  specifications by reducing it to realizability subproblems for safety specifications. The reduction to safety, which we will give in Sec. 5, generates a safety formula for each integer  $k$ , in such a way to guarantee the following important properties: (i) *soundness*, ensuring that the realizability of the  $k^{\text{th}}$  subproblem implies the realizability of the starting formula, and (ii) *completeness*, establishing the existence of an upper bound  $\mu$  such that the unrealizability of all the  $k^{\text{th}}$  subproblems with  $k \leq \mu$  implies the unrealizability of the starting formula. In the next section we will give a general framework for (sound and) complete reductions. From it, in Sec. 5, we will derive one for  $\text{GR-EBR}$  specifications, showing also how the realizability of each safety subproblems can be solved symbolically.

## 4 A Framework of Safety Reductions for $\text{LTL}+\text{P}$ Realizability

The central question of this section is: *how can we obtain a complete safety reduction for the realizability problem of specifications written in (fragments of) LTL?* In the following, we propose a framework to answer it.

The core and the main novelty of our framework is a link with safety reductions for model checking: in order to design a complete reduction for the realizability problem, one can prove that it is complete for the model checking problem and then use our framework to derive completeness for realizability. On one hand, this allows to prove completeness at the level of model checking, which is simpler than proving completeness for realizability. On the other hand, this opens the possibility of using existing safety reductions already devised for model checking for realizability as well. We start by defining what is a safety reduction in the context of our framework.

**Definition 9 (Safety reduction).** *Let  $\mathcal{S} \subseteq \text{LTL}$  be a fragment of LTL. A safety reduction for  $\mathcal{S}$  is a function  $\llbracket \cdot \rrbracket$  such that, for each formula  $\phi \in \mathcal{S}$  over*



the alphabet  $\Sigma$ , it holds that  $\llbracket \phi \rrbracket = \{\phi_k\}_{k \in \mathbb{N}}$ , where  $\phi_k$  is a safety formula over the alphabet  $\Sigma$  such that  $\phi_k \rightarrow \phi$ , for any  $k \in \mathbb{N}$ . With  $\llbracket \phi \rrbracket^k$ , we will denote the formula  $\phi_k$  of the set above.

*Link between realizability and model checking* The rationale behind the link between realizability and model checking is the following one: since we can easily view Mealy machines as (a particular type of) Kripke structures and viceversa, and since by Prop. 1 we can restrict realizability to the search of finite strategies representable by Mealy machines, the realizability problem of the LTL+P formula  $\phi$  can be reduced to checking if there exists a Mealy machine  $M_g$  such that  $M'_g \models A\phi$ , where  $M'_g$  is the Kripke structure *corresponding to*  $M_g$ .

The Kripke structure  $M'_g$  *corresponding to* the Mealy machine  $M_g = (2^{\mathcal{U}}, 2^{\mathcal{C}}, Q, q_0, \delta)$  is defined as  $M'_g = (2^{\mathcal{U} \cup \mathcal{C}}, Q', I', T', L')$  where:

1.  $Q' = Q \times \{q_U \mid U \in 2^{\mathcal{U}}\} \times \{q_C \mid C \in 2^{\mathcal{C}}\}$ ;
2.  $I' = \{(q_0, q_U, q_C) \in Q' \mid \delta(q_0, U) = (C, q') \text{ for any } U \in 2^{\mathcal{U}}, C \in 2^{\mathcal{C}} \text{ and } q' \in Q\}$ ,
3.  $T' = \{((q, q_U, q_C), (q', q_{U'}, q_{C'})) \mid \delta(q, U) = (C, q') \text{ for any } U, U' \in 2^{\mathcal{U}}, C, C' \in 2^{\mathcal{C}}, \text{ and } q, q' \in Q\}$  and
4.  $L'((q, q_U, q_C)) = U \cup C$ .

The Kripke structure  $M'_g$  is such that each trace of  $M'_g$  corresponds to a word of  $M_g$ , and viceversa.

In proving the completeness theorem, we will abstract from the concrete safety reduction and give the conditions for a general safety reduction  $\llbracket \cdot \rrbracket$  (as defined in Def. 9) to be complete. These conditions are formalized in Def. 10.

**Definition 10 (Sound and Complete safety reduction).** *Let  $\mathcal{S} \subseteq \text{LTL}$  be a fragment of LTL,  $\phi$  a formula in  $\mathcal{S}$ , and  $\llbracket \cdot \rrbracket$  a safety reduction for  $\mathcal{S}$ . We say that  $\llbracket \cdot \rrbracket$  is  $\mu$ -complete, for a given function  $\mu : \mathbb{N} \rightarrow \mathbb{N}$  if and only if, for all  $\phi \in \mathcal{S}$  and for all Kripke structures  $M$ :*

$$M \models A\phi \quad \Leftrightarrow \quad \exists k \leq \mu(|M|) . M \models A\llbracket \phi \rrbracket^k$$

We can finally state the main theorem of our framework, which uses Def. 10 and Prop. 1 in order to establish that if a safety reduction is complete for the model checking problem, then it is complete for the realizability problem as well.

**Theorem 1 (Soundness and Completeness for LTL+P Realizability).** *Let  $\mathcal{S} \subseteq \text{LTL}$  be a fragment of LTL,  $\phi \in \mathcal{S}$  a formula over the input alphabet  $\mathcal{U}$  and output alphabet  $\mathcal{C}$  (with  $n = |\phi|$ ) and  $\llbracket \cdot \rrbracket$  a  $\mu$ -complete safety reduction for  $\mathcal{S}$ , for a given function  $\mu$ . It holds that:*

$$\phi \text{ is realizable} \quad \Leftrightarrow \quad \exists k \leq \mu(2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot 2^{2^{c \cdot n}}) . \llbracket \phi \rrbracket^k \text{ is realizable}$$

*Proof.* We first prove the *soundness*, which corresponds to the right-to-left direction. Suppose there exist a  $k \leq \mu(2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot 2^{2^{c \cdot n}})$  such that  $\llbracket \phi \rrbracket^k$  is realizable.

Then, there exists a strategy  $g : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$  such that  $\mathcal{L}(g) \subseteq \mathcal{L}(\llbracket \phi \rrbracket^k)$ . By Prop. 1, there exists a Mealy machine  $M_g = (2^{\mathcal{U}}, 2^{\mathcal{C}}, Q, q_0, \delta)$  with input alphabet  $2^{\mathcal{U}}$  and output alphabet  $2^{\mathcal{C}}$  such that  $\mathcal{L}(M_g) \subseteq \mathcal{L}(\llbracket \phi \rrbracket^k)$ . Starting from  $M_g$ , let  $M'_g = (2^{\mathcal{U} \cup \mathcal{C}}, Q', I', T', L')$  be the *corresponding* Kripke structure. The Kripke structure  $M'_g$  is such that each trace of  $M'_g$  corresponds to a word of  $M_g$ , and viceversa. Therefore all the traces  $\pi$  of  $M'_g$  are such that  $L'(\pi) \models \llbracket \phi \rrbracket^k$ , that is  $M'_g \models \mathbf{A}\llbracket \phi \rrbracket^k$ . Since by hypothesis  $\llbracket \cdot \rrbracket$  is a  $\mu$ -complete safety reduction, by Def. 10, it holds that  $M'_g \models \mathbf{A}\phi$ . This means that also  $\mathcal{L}(M_g) \subseteq \mathcal{L}(\phi)$ . Since  $M_g$  is a Mealy machine, this implies that  $\phi$  is realizable.

We now prove *completeness*, which corresponds to the left-to-right direction. Suppose that  $\phi$  is realizable. Since  $\phi \in S$  and since  $S \subseteq \text{LTL}$ ,  $\phi$  is an LTL formula as well. Therefore, by Prop. 1, there exists a Mealy machine  $M_g$  with input alphabet  $2^{\mathcal{U}}$  and output alphabet  $2^{\mathcal{C}}$  such that  $\mathcal{L}(M_g) \subseteq \mathcal{L}(\phi)$  with at most  $2^{2^{c \cdot n}}$  states, for some constant  $c \in \mathbb{N}$ . From  $M_g$ , we build an equivalent Kripke structure  $M'_g$  with input alphabet  $\Sigma' = 2^{\mathcal{U} \cup \mathcal{C}}$ , as described above for the soundness proof. It holds that  $M'_g \models \mathbf{A}\phi$ . Since by hypothesis  $\llbracket \cdot \rrbracket$  is a  $\mu$ -complete safety reduction for  $S$ , and since  $|Q'| = 2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot |Q|$  (where  $Q$  and  $Q'$  are the set of states of  $M_g$  and  $M'_g$ , respectively), by Def. 10, there exists a  $k \leq \mu(2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot 2^{2^{c \cdot n}})$  such that  $M'_g \models \mathbf{A}\llbracket \phi \rrbracket^k$ . This means that also  $\mathcal{L}(M_g) \subseteq \mathcal{L}(\llbracket \phi \rrbracket^k)$ . Since  $M_g$  is a Mealy machine, this means that there exists a  $k \leq \mu(2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot 2^{2^{c \cdot n}})$  such that  $\llbracket \phi \rrbracket^k$  is realizable.  $\square$

*Novelty and Usage* As already mentioned before, a distinguished and important feature of our framework is that it provides a link with safety reductions for the *model checking problem*. This opens the possibility to use model checking safety reductions for the realizability problem as well, provided that the reduction fulfills the requirements in Def. 10. In the next sections, we will define a *concrete* safety reduction for GR-EBR specifications that is *complete* with respect to Def. 10, and we will use it for introducing a novel algorithm for GR-EBR realizability. Using Theorem 1, we will derive a corollary for the completeness of our algorithm.

*In practice* The upper bound for the value of  $\mu(\cdot)$  (after which we can answer *unrealizable*) is doubly exponential in the size of the initial formula and therefore, in practice, it is prohibitively large. It follows that usually the *completeness* of a safety reduction can be exploited in practice only for making sure that, starting from a *realizable* specification, we will eventually find a  $k \in \mathbb{N}$  such that the  $k^{\text{th}}$  subproblem is realizable. Therefore, like K-Liveness for model checking [6], we can use our algorithm in parallel with another one that checks for the unrealizability of the specification. The first that terminates stops the other and, thus, the entire procedure. We remark that we cannot check the unrealizability of  $\phi$  by solving the dualized game (*i.e.*, looking for a Moore-type strategy of Environment) for  $\neg\phi$ , because GR-EBR and  $\text{LTL}_{\text{EBR}}+\text{P}$  are *not* closed under complementation.

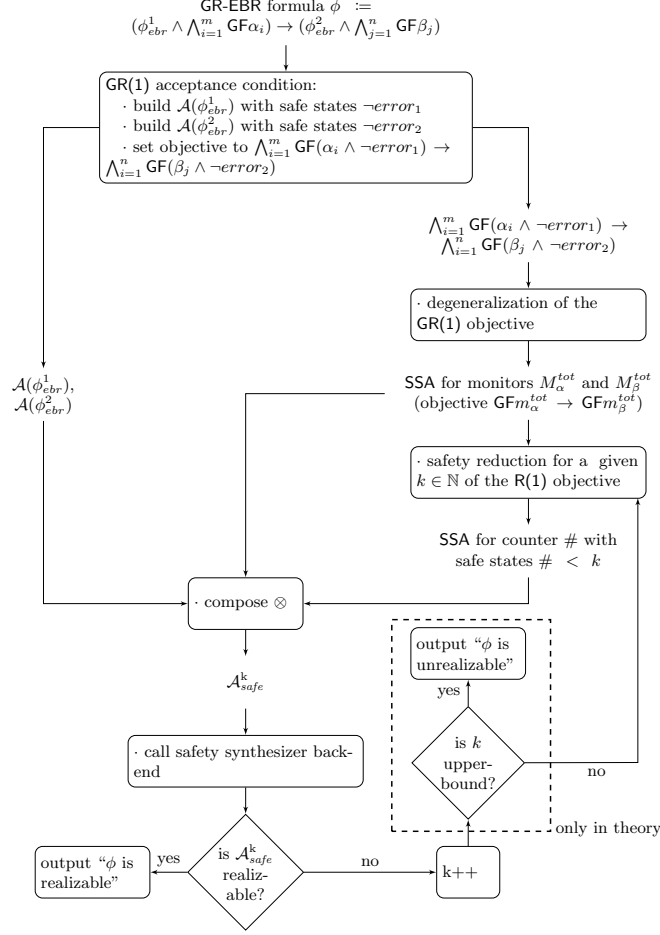


Fig. 1: Low-level view of the procedure for the realizability of GR-EBR formulas.

## 5 A Safety Reduction for GR-EBR

In this section, we describe the algorithm for solving realizability of GR-EBR specifications. It consists in three steps. Firstly, we build the product between the two symbolic and safety automata for the safety parts of both assumptions and guarantees. This product automaton has a GR(1) accepting condition. The second step consists in a so-called *degeneralization*, that, by using deterministic monitors, turns the GR(1) accepting condition into a Reactivity(1) (R(1), for short) condition. The third and last step, that is the core of the procedure, reduces the realizability problem over the above automaton to a *sequence* of safety synthesis problems, that is, realizability problems over safety (and symbolic) automata  $\mathcal{A}_{safe}^k$ , one for each index  $k \in \mathbb{N}$ . By introducing a *concrete* safety reduction  $\llbracket \cdot \rrbracket_{ebr}$  for GR-EBR, and by proving that it is *complete* with respect to

Def. 10, we prove the completeness of the entire procedure. The structure of the full procedure is depicted in Fig. 1.

Finally, note that, as for now, there is no incrementality between an iteration and the next one, because of the lack of incremental safety synthesizers. The only point that we save between one iteration and the next one is the construction of the two symbolic safety automata, which is performed only once during the procedure.

### 5.1 Construction of the automaton with a GR(1) condition

In this part, we describe the first step of the algorithm. Starting from a GR-EBR formula  $\phi := (\phi_{ebr}^1 \wedge \bigwedge_{i=1}^m \text{GF}\alpha_i) \rightarrow (\phi_{ebr}^2 \wedge \bigwedge_{j=1}^n \text{GF}\beta_j)$ , the objective is to obtain an automaton  $\mathcal{A}$  such that: (i) it has a GR(1) accepting condition, and (ii) it recognizes the same language of  $\phi$ , *i.e.*,  $\mathcal{L}(\phi) = \mathcal{L}(\mathcal{A})$ . In order to do that, we first build the two symbolic safety automata for the safety parts of both the assumptions and the guarantees, that is for  $\phi_{ebr}^1$  and  $\phi_{ebr}^2$ . Since by definition both are  $\text{LTL}_{\text{EBR}}+\text{P}$  formulas, we use the transformation described in [4], to which the reader is referred for more details.

From now on, let  $\mathcal{A}(\phi_{ebr}^1)$  and  $\mathcal{A}(\phi_{ebr}^2)$  be the automata for  $\phi_{ebr}^1$  and  $\phi_{ebr}^2$ , respectively. Let  $\mathcal{A}_{\phi_{ebr}}$  be the product automaton  $\mathcal{A}(\phi_{ebr}^1) \times \mathcal{A}(\phi_{ebr}^2)$ . The question is how to set the acceptance condition of  $\mathcal{A}_{\phi_{ebr}}$  such that the conditions (i) and (ii) of above are fulfilled. We answer this question by examining how the automata  $\mathcal{A}(\phi_{ebr}^1)$  and  $\mathcal{A}(\phi_{ebr}^2)$  are made internally. Take for example the formula  $\text{G}p$  (for some atomic proposition  $p \in \Sigma$ ). The safety automaton corresponding to this formula comprises an *error bit* as one of its state variables, let us call it **error**, which is initially set to be **false**. The transition function for **error** is deterministic and updates **error** to **true** if  $\neg p$  holds in the current state, or keeps its value otherwise. The set of safe states comprises all and only those states in which **error** is **false**. In a symbolic setting, this is expressed by the formula  $\text{G}\neg\text{error}$ . In this way,  $p$  is forced to hold constantly in all (and only) the words accepted by the automaton.

A crucial property of each error bit is *monotonicity*: once **error** is set to **true**, it can never be set to **false** again. Formally, given a trace  $\tau$  of the automaton, it holds that, if there exists  $i \geq 0$  such that  $\tau(i) \models \text{error}$ , then  $\tau(j) \models \text{error}$ , for all  $j \geq i$ . Monotonicity of the error bits allows us to express an accepting condition of type  $\text{G}\neg\text{error}$  in terms of  $\text{GF}\neg\text{error}$ , by maintaining the equivalence.

**Lemma 1 (Monotonicity of Error Bits).** *Each error bit is monotone.*

*Proof.* Consider a trace  $\tau$  of an automaton with an accepting condition of the type  $\text{G}\neg\text{error}$ . If  $\tau \models \text{G}\neg\text{error}$  then of course  $\tau \models \text{GF}\neg\text{error}$ . Suppose now that  $\tau \models \text{GF}\neg\text{error}$ . If by contradiction we suppose that  $\tau \not\models \text{G}\neg\text{error}$ , we have that there exists an  $i \geq 0$  such that  $\tau(i) \models \text{error}$ . By the monotonicity property, this would mean that also  $\tau(j) \models \text{error}$ , for all  $j \geq i$ , that is  $\tau \models \text{FGerror}$ , but this a contradiction with our hypothesis. Therefore, we proved that changing the acceptance condition of an automaton from a  $\text{G}\neg\text{error}$  to  $\text{GF}\neg\text{error}$  maintains the equivalence.  $\square$

Let  $\mathbf{error}_1$  and  $\mathbf{error}_2$  be the error bits of  $\mathcal{A}(\phi_{ebr}^1)$  and  $\mathcal{A}(\phi_{ebr}^2)$ , respectively. Let  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$  be the automaton obtained from  $\mathcal{A}_{\phi_{ebr}}$  by replacing its acceptance condition with the following GR(1) condition:

$$(\mathbf{GF}\neg\mathbf{error}_1 \wedge \bigwedge_{i=1}^m \mathbf{GF}\alpha_i) \rightarrow (\mathbf{GF}\neg\mathbf{error}_2 \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j) \quad (1)$$

The intuition is that  $\mathbf{error}_1$  and  $\mathbf{error}_2$  keep track of the *safety* parts of  $\phi$ , that is  $\phi_{ebr}^1$  and  $\phi_{ebr}^2$ . The following lemma proves the equivalence between  $\phi$  and  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$ .

**Lemma 2.** *Let  $\phi$  be an GR-EBR formula. It holds that  $\mathcal{L}(\phi) = \mathcal{L}(\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge})$ .*

*Proof.* Let  $\phi \in \text{GR-EBR}$ .  $\phi$  is of the following form:

$$(\phi_{ebr}^1 \rightarrow \bigwedge_{i=1}^m \mathbf{GF}\alpha_i) \rightarrow (\phi_{ebr}^2 \rightarrow \bigwedge_{j=1}^n \mathbf{GF}\beta_j)$$

By the theorems proved in [4], it holds that  $\mathcal{L}(\phi_{ebr}^1) = \mathcal{L}(\mathcal{A}(\phi_{ebr}^1))$  and  $\mathcal{L}(\phi_{ebr}^2) = \mathcal{L}(\mathcal{A}(\phi_{ebr}^2))$ .

Consider first the left-to-right direction. Let  $\sigma \in \mathcal{L}(\phi)$ . We prove that  $\sigma \in \mathcal{L}(\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge})$ . Each  $\sigma \in \mathcal{L}(\phi)$  is such that: a. either  $\sigma \models \neg\phi_{ebr}^1 \vee \neg(\bigwedge_{i=1}^m \mathbf{GF}\alpha_i)$ , b. or  $\sigma \models \phi_{ebr}^1 \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ . Recall that  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$  is defined as the product automaton  $\mathcal{A}(\phi_{ebr}^1) \times \mathcal{A}(\phi_{ebr}^2)$  with the acceptance condition  $\alpha$  defined as  $(\mathbf{GF}\neg\mathbf{error}_1 \wedge \bigwedge_{i=1}^m \mathbf{GF}\alpha_i) \rightarrow (\mathbf{GF}\neg\mathbf{error}_2 \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j)$ .

Consider case *a*. If  $\sigma \models \neg\phi_{ebr}^1 \vee \neg(\bigwedge_{i=1}^m \mathbf{GF}\alpha_i)$ , then the trace induced by  $\sigma$  in  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$  is such that at least one of the following two cases hold:

- a.1. either  $\exists i \geq 0$  such that  $\tau(i) \models \mathbf{error}_1$ , that is  $\tau \models \mathbf{F}(\mathbf{error}_1)$ . In this case, we exploit *monotonicity* of  $\mathbf{error}_1$ . Since  $\tau \models \mathbf{F}(\mathbf{error}_1)$ , it also holds that  $\tau \models \mathbf{FG}(\mathbf{error}_1)$ , that is  $\tau \not\models \mathbf{GF}(\neg\mathbf{error}_1)$ . As a consequence,  $\tau \models \alpha$ , where  $\alpha$  is the acceptance condition of  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$ , and thus  $\sigma \in \mathcal{L}(\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge})$ .
- a.2. or  $\tau \models \neg\bigwedge_{i=1}^m \mathbf{GF}\alpha_i$ . In this case, of course,  $\tau \models \alpha$  (that is,  $\tau$  satisfies the acceptance condition of  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$ ), and thus  $\sigma \in \mathcal{L}(\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge})$ .

Consider now the case *b*. If  $\sigma \models \phi_{ebr}^1 \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ , then  $\sigma \models \phi_{ebr}^1$  and  $\sigma \models \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ . Therefore, the trace induced by  $\sigma$  in  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$  is such that  $\tau \models \mathbf{G}(\neg\mathbf{error}_2) \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ , that implies that  $\tau \models \mathbf{GF}(\neg\mathbf{error}_2) \wedge \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ . Therefore,  $\tau \models \alpha$ , and thus  $\sigma \in \mathcal{L}(\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge})$ . The opposite direction can be proved similarly.  $\square$

## 5.2 Degeneralization

The objective of this part is to transform the GR(1) accepting condition of the automaton  $\mathcal{A}_{ebr}^{\wedge \rightarrow \wedge}$ , that is of the form  $\bigwedge_{i=1}^m \mathbf{GF}\alpha_i \rightarrow \bigwedge_{j=1}^n \mathbf{GF}\beta_j$ , into a condition of the form  $\mathbf{GF}\alpha \rightarrow \mathbf{GF}\beta$  (also called *Reactivity(1)* objective,  $\mathbf{R}(1)$ , for short). In this context, we will use the term *monitor* as a synonym of *deterministic*

*automaton*. In order to accomplish the task, for each  $\alpha_i$  (resp. for each  $\beta_i$ ), we define a monitor  $M_{\alpha_i}$  (resp.  $M_{\beta_i}$ ) that is set to *true* when  $\alpha_i$  (resp.  $\beta_i$ ) has been read and is reset to *false* when all the  $\alpha_i$  (resp.  $\beta_i$ ) have been read. For this last condition, we define the monitors  $M_{\alpha}^{tot}$  and  $M_{\beta}^{tot}$ .

Let  $M_{\alpha_i}$  and  $M_{\alpha}^{tot}$  be the symbolic safety automata such that their input alphabet is  $2^{\Sigma}$  (where  $\Sigma$  is the alphabet of the starting GR-EBR formula), their set of *state variables* are  $\{m_{\alpha_i}\}$  and  $\{m_{\alpha}^{tot}\}$ , respectively, all their reachable states are safe states, and their transition relations are the following:

|  |   |
|--|---|
| $\text{init}(m_{\alpha_i}) := 0$           | $\text{init}(m_{\alpha}^{tot}) := 0$                |
| $\text{next}(m_{\alpha_i}) := \text{case}$ | $\text{next}(m_{\alpha}^{tot}) := \text{case}$      |
| $\alpha_i : 1$                             | $m_{\alpha_1} \wedge \dots \wedge m_{\alpha_m} : 1$ |
| $m_{\alpha}^{tot} : 0$                     | $\text{default} : 0$                                |
| $\text{default} : m_{\alpha_i}$            | $\text{esac}$                                       |
| $\text{esac}$                              |   |

We define  $M_{\beta_i}$  and  $M_{\beta}^{tot}$  as  $M_{\alpha_i}$  and  $M_{\alpha}^{tot}$ , respectively, but with  $\alpha_i$  substituted with  $\beta_i$  and  $\alpha$  substituted with  $\beta$ . Let  $\mathcal{A}_{degen}$  be the product between all the  $M_{\alpha_i}$ ,  $M_{\beta_i}$ ,  $M_{\alpha}^{tot}$  and  $M_{\beta}^{tot}$ . Let  $\mathcal{A}_{degen}^{\text{GF} \rightarrow \text{GF}}$  be the automaton obtained from  $\mathcal{A}_{degen}$  by replacing its accepting condition with the Reactivity(1) condition  $\text{GF}m_{\alpha}^{tot} \rightarrow \text{GF}m_{\beta}^{tot}$ . We can prove the following lemma, which states that this step of the algorithm maintains the equivalence.

**Lemma 3.**  $\mathcal{L}(\mathcal{A}_{\text{ebr}}^{\wedge \rightarrow \wedge}) = \mathcal{L}(\mathcal{A}_{\phi_{\text{ebr}}} \times \mathcal{A}_{\text{degen}}^{\text{GF} \rightarrow \text{GF}})$ .

*Proof.* We prove separately the two directions. Consider first the right-to-left direction. Let  $\sigma$  be an infinite word of  $\mathcal{L}(\mathcal{A}_{\phi_{\text{ebr}}} \times \mathcal{A}_{\text{degen}}^{\text{GF} \rightarrow \text{GF}})$ . Then  $\sigma$  is a word in  $\mathcal{L}(\mathcal{A}_{\phi_{\text{ebr}}})$ . Moreover,  $\sigma$  is a word in  $\mathcal{L}(\mathcal{A}_{\text{degen}}^{\text{GF} \rightarrow \text{GF}})$  and thus there exists a run  $\tau$  induced by  $\sigma$  such that  $\tau \models \text{GF}m_{\alpha}^{tot} \rightarrow \text{GF}m_{\beta}^{tot}$ , that is,  $\tau \models \text{FG}\neg m_{\alpha}^{tot} \vee \text{GF}m_{\beta}^{tot}$ . We divide in cases:

- if  $\tau \models \text{FG}\neg m_{\alpha}^{tot}$ , then by the semantics of the temporal operators F and G, there exists an  $i \geq 0$  such that for all  $j \geq i$ ,  $\tau_j \models \neg m_{\alpha}^{tot}$ . By construction of the monitors  $m_{\alpha}^{tot}$ , this means that there exists an  $i \geq 0$  such that for all  $j \geq i$ ,  $\tau_j \models \bigvee_{k=1}^m \neg m_{\alpha_k}$ . This implies that, there exists a  $k \in [1, m]$  and an  $i \geq 0$  such that for all  $j \geq i$ , such that  $\tau_j \models \neg m_{\alpha_k}$ . Indeed, suppose by contradiction that it is not so: then for all  $k \in [1, m]$ , there exists infinitely many positions  $i \geq 0$  such that  $\tau_i \models m_{\alpha_k}$ . This would mean that the monitor  $M_{\alpha}^{tot}$  is set to *true* infinitely many times, that is  $\text{GF}m_{\alpha}^{tot}$ , but this is a contradiction with our hypothesis. Therefore, it holds that  $\tau \models \bigvee_{k=1}^m \text{FG}\neg m_{\alpha_k}$ , and thus also that  $\tau \models \bigwedge_{i=1}^m \text{GF}\alpha_i \rightarrow \bigwedge_{j=1}^n \text{GF}\beta_j$ . Overall, since  $\tau$  is induced by  $\sigma$ , we have that  $\sigma$  is a word of  $\mathcal{L}(\mathcal{A}_{\phi_{\text{ebr}}})$  that induces a run  $\tau$  such that  $\tau \models \bigwedge_{i=1}^m \text{GF}\alpha_i \rightarrow \bigwedge_{j=1}^n \text{GF}\beta_j$ , that is  $\sigma \in \mathcal{L}(\mathcal{A}_{\text{ebr}}^{\wedge \rightarrow \wedge})$ .
- If otherwise  $\tau \models \text{GF}m_{\beta}^{tot}$ , then there exists infinitely many positions  $i \geq 0$  such that  $\tau_i \models m_{\beta}^{tot}$ . Moreover, it holds that for all  $i_1 \geq 0$  and for all  $i_2 \geq i_1$ , if  $\tau_{i_1} \models m_{\beta}^{tot}$  and  $\tau_{i_2} \models m_{\beta}^{tot}$ , then, for all  $1 \leq k \leq n$ , there exists a  $i_1 \leq j \leq i_2$  such that  $\tau_j \models m_{b_k}$ . Putting together these two points, we have that for all  $1 \leq k \leq n$ , there exists infinitely many  $i \geq 0$  such that  $\tau_i \models m_{b_k}$ . That is,  $\tau \models \bigwedge_{k=1}^n \text{GF}m_{b_k}$ . By definition of the monitors  $M_{\beta_i}$  and since  $\tau$  is induced

by  $\sigma$ , we have that  $\sigma$  is a word in  $\mathcal{L}(\mathcal{A}_{\phi_{\text{ebr}}})$  that induces a run  $\tau$  such that  $\tau \models \bigwedge_{i=1}^m \text{GF}\alpha_i \rightarrow \bigwedge_{j=1}^n \text{GF}\beta_j$ . That is,  $\sigma \in \mathcal{L}(\mathcal{A}_{\text{ebr}}^{\wedge \rightarrow \wedge})$ .

The proof the left-to-right direction is specular, and therefore is omitted from the presentation.  $\square$

### 5.3 Reduction to Safety for Reactivity(1) objectives

In this part, we describe a *complete* safety reduction (see Def. 10) tailored for Reactivity(1) objectives. We will apply this reduction on the automaton  $\mathcal{A}_{\text{degen}}^{\text{GF} \rightarrow \text{GF}}$  obtained from the previous step. The intuition is to use a *counter* to count and limit the number of positions, after a position in which  $m_{\beta}^{\text{tot}}$  holds, in which  $m_{\alpha}^{\text{tot}} \wedge \neg m_{\beta}^{\text{tot}}$  holds. We define the counter as follows.

**Definition 11 (Counter for the Reactivity(1) objective).** Let  $\mathcal{A}_{\#_{\alpha,\beta}^{\rightarrow}}^k$  be the symbolic and deterministic safety automaton whose set of safe states is represented by the formula  $\text{G}(\#_{\alpha,\beta}^{\rightarrow} < k)$  and whose transition relation is the following:

---

|   |    |                                       |
|---|----|---------------------------------------|
| $init(\#_{\alpha,\beta}^{\rightarrow})$ | := | $0$                                   |
| $next(\#_{\alpha,\beta}^{\rightarrow})$ | := | <i>case</i>                           |
| $m_{\beta}^{\text{tot}}$                | :  | $0$                                   |
| $m_{\alpha}^{\text{tot}}$               | :  | $\#_{\alpha,\beta}^{\rightarrow} + 1$ |
| <i>default</i>                          | :  | $\#_{\alpha,\beta}^{\rightarrow}$     |

---

We define  $\mathcal{A}_{\text{safe}}^k := \mathcal{A}_{\phi_{\text{ebr}}} \times \mathcal{A}_{\text{degen}} \times \mathcal{A}_{\#_{\alpha,\beta}^{\rightarrow}}^k$ , and we set the accepting condition of  $\mathcal{A}_{\text{safe}}^k$  to be the one of  $\mathcal{A}_{\#_{\alpha,\beta}^{\rightarrow}}^k$ , i.e.,  $\text{G}(\#_{\alpha,\beta}^{\rightarrow} \leq \#_{\alpha,\beta}^{\rightarrow} < k)$ . The automaton  $\mathcal{A}_{\text{safe}}^k$  is a symbolic and deterministic safety automaton, and therefore it can be used as an arena for a safety game. In practice, we check the realizability of  $\mathcal{A}_{\text{safe}}^k$  by means of a tool for safety synthesis. We start with  $k = 0$ , and we check the realizability of  $\mathcal{A}_{\text{safe}}^k$ : if Controller has a strategy, then we stop, otherwise we increment  $k$  and we repeat the cycle.

In order to prove that this step is sound and complete, we use the framework described in Sec. 4. We call  $\llbracket \cdot \rrbracket_{\text{ebr}}$  the safety reduction described in this part. Since the framework works with formulas rather than with automata, for all  $\phi \in \text{GR-EBR}$ , we define  $\llbracket \phi \rrbracket_{\text{ebr}}^k$  to be any *safety formula* such that  $\mathcal{L}(\llbracket \phi \rrbracket_{\text{ebr}}^k) = \mathcal{L}(\mathcal{A}_{\text{safe}}^k)$ . From now, with  $\text{id} : \mathbb{N} \rightarrow \mathbb{N}$  we denote the *identity* function.

**Theorem 2.**  $\llbracket \cdot \rrbracket_{\text{ebr}}$  is a *id-complete safety reduction* for *GR-EBR*.

*Proof.* We have to prove that, for all  $\phi \in \text{GR-EBR}$ , for all Kripke structures  $M$  and for all  $k \in \mathbb{N}$ , it holds that:

$$M \models \text{A}\phi \quad \Leftrightarrow \quad \exists k \leq \text{id}(|M|) . M \models \text{A}\llbracket \phi \rrbracket_{\text{ebr}}^k$$

We prove separately the two directions. Consider first the *soundness* which corresponds to the right-to-left direction. Suppose that  $M \models \text{A}\llbracket \phi \rrbracket_{\text{ebr}}^k$ . It holds

that, for each initialized trace  $\pi$  of  $M$ ,  $L(\pi) \models \llbracket \phi \rrbracket_{ebr}^k$ , where  $L(\cdot)$  is the labeling function of  $M$ . Let  $\pi$  be an initialized trace of  $M$ . By definition of  $\llbracket \cdot \rrbracket_{ebr}$ , it holds that, there exists a run  $\tau$  induced by  $L(\pi)$  such that: (i)  $\tau$  is accepting in  $\mathcal{A}_{\phi_{ebr}} \times \mathcal{A}_{degen}$ , and (ii)  $\tau$  is accepting in  $\mathcal{A}_{\#_{\alpha,\beta}^k}$ . From the second point, we have that:

- either,  $\#_{\alpha,\beta}^{\rightarrow}$  make infinitely many *resets*. This means that there exists infinitely many positions in  $\tau$  in which  $m_{\alpha}^{tot}$  holds and, after at most  $k$  occurrences of  $m_{\alpha}^{tot}$ , there is a  $m_{\beta}^{tot}$ . Therefore, in particular, there exists infinitely many positions in which  $m_{\beta}^{tot}$  holds, that is  $\tau \models \text{GF}m_{\beta}^{tot}$ .
- or the counter  $\#_{\alpha,\beta}^{\rightarrow}$  stops to increment because, because it does not read any  $m_{\alpha}^{tot}$ . This means that there exists *finitely* many positions in which  $m_{\alpha}^{tot}$  holds, that is  $\tau \models \text{FG}\neg m_{\alpha}^{tot}$ .

Therefore, it holds that  $\tau \models \text{FG}\neg m_{\alpha}^{tot} \vee \text{GF}m_{\beta}^{tot}$ , that is  $\tau \models \text{GF}m_{\alpha}^{tot} \rightarrow \text{GF}m_{\beta}^{tot}$ . Finally, we have that  $\tau$  is an accepting run of  $\mathcal{A}_{\phi_{ebr}} \times \mathcal{A}_{degen}$  such that  $\tau \models \text{GF}m_{\alpha}^{tot} \rightarrow \text{GF}m_{\beta}^{tot}$ . Since by hypothesis  $L(\pi)$  is induced by  $\tau$ , by definition of  $\mathcal{A}_{degen}^{\text{GF} \rightarrow \text{GF}}$ , we have that  $L(\pi) \in \mathcal{L}(\mathcal{A}_{\phi_{ebr}} \times \mathcal{A}_{degen}^{\text{GF} \rightarrow \text{GF}})$ . By concatenating Lemma 2 and Lemma 3, we have that  $L(\pi) \in \mathcal{L}(\phi)$ , and therefore  $\pi \models \phi$ . It follows that  $M \models \text{A}\phi$ .

We now prove *completeness*, which corresponds to the left-to-right direction. Suppose that  $M \models \text{A}\phi$ , where  $\phi \in \text{GR-EBR}$ . We prove this case by contradiction. Suppose therefore that for all  $k \leq \text{id}(|M|)$ ,  $M \not\models \text{A}\llbracket \phi \rrbracket_{ebr}^k$ . This means that there exists an initialized trace  $\pi$  in  $M$  such that  $L(\pi) \notin \mathcal{L}(\llbracket \phi \rrbracket_{ebr}^k)$ , for all  $k \leq \text{id}(|M|)$ . By definition of  $\llbracket \cdot \rrbracket_{ebr}$ , for  $k = \text{id}(|M|)$ , we have that *for all* runs  $\tau$  induced by  $L(\pi)$  in  $\mathcal{A}_{\phi_{ebr}} \times \mathcal{A}_{degen} \times \mathcal{A}_{\#_{\alpha,\beta}^k}$ , it holds that  $\tau \not\models \text{G}(\#_{\alpha,\beta}^{\rightarrow} \leq k)$ . Let  $\tau$  be one of these runs. There exists a position  $i$  in  $\tau$  such that  $\tau_i \models (\#_{\alpha,\beta}^{\rightarrow} = v)$ , for some  $v > k$ . By definition of the counter  $\#_{\alpha,\beta}^{\rightarrow}$ , the run  $\tau$  is such that:

$$\begin{aligned} \exists 0 < h_1 < h_2 < \dots < h_v . ( \tau_{h_1} \models m_{\alpha}^{tot} \wedge \tau_{h_2} \models m_{\alpha}^{tot} \wedge \dots \wedge \tau_{h_v} \models m_{\alpha}^{tot} \wedge \\ \forall h_1 \leq h \leq h_v . (\tau_h \models \neg m_{\beta}^{tot})) \end{aligned}$$

Recall that  $\tau$  is a run induced by  $L(\pi)$ . Since  $v > k$ ,  $k = \text{id}(|M|)$  and  $M$  is a *finite-state* Kripke structure, the positions  $h_1 \dots h_v$  in  $\pi$  (attention: *not* in  $\tau$ ) cannot be all different. That is, there exists at least two indexes  $s, e \in \mathbb{N}$  such that: (i)  $1 \leq s < e \leq v$ , (ii)  $\pi_{h_s} = \pi_{h_e}$ , and (iii)  $\pi_{h_s} \models m_{\alpha}^{tot}$ . Starting from  $\pi$ , we can build a *looping trace*  $\pi'$  that agrees with  $\pi$  in the prefix  $\pi_{[0, h_e]}$  and then loops on the interval  $\pi_{[h_s, h_e]}$ . It holds that  $\pi'$  is an initialized trace of  $M$  and it induces a run  $\tau'$  such that  $\tau' \models \text{GF}m_{\alpha}^{tot} \wedge \text{FG}\neg m_{\beta}^{tot}$ , that is  $\tau' \not\models \text{GF}m_{\alpha}^{tot} \rightarrow \text{GF}m_{\beta}^{tot}$ . Nevertheless, since  $M \models \text{A}\phi$ , by Lemma 2 and Lemma 3, we have that  $L(\pi') \in \mathcal{L}(\mathcal{A}_{\phi_{ebr}} \times \mathcal{A}_{degen}^{\text{GF} \rightarrow \text{GF}})$ , and therefore this is a contradiction. This means that it has to hold that  $L(\pi) \in \mathcal{L}(\llbracket \phi \rrbracket_{ebr}^k)$ , that is  $\pi \models \llbracket \phi \rrbracket_{ebr}^k$  for all the initialized traces  $\pi$  of  $M$ , and thus there exists a  $k \leq \text{id}(|M|)$  such that  $M \models \text{A}\llbracket \phi \rrbracket_{ebr}^k$ .  $\square$

With Theorem 1, we derive the following corollary that proves the completeness of our procedure.



**Corollary 1.** *For any formula  $\phi \in GR\text{-EBR}$ , it holds that:  $\phi$  is realizable iff  $\exists k \leq \text{id}(2^{|\mathcal{U}|} \cdot 2^{|\mathcal{C}|} \cdot 2^{2^{c-n}})$  such that  $\llbracket \phi \rrbracket_{\text{ebr}}^k$  is realizable.*

## 6 Experimental Evaluation

We implemented the algorithm described in Sec. 5 and summarized in Fig. 1 in a prototype tool called GRACE (which stands for *GR-ebr reAlizability ChEcker*)<sup>4</sup>. We chose SAFETYSYNTH [14] as a BDD-based backend for solving each safety game.

As competitor tools, we chose BOsY [9, 10, 12] and STRIX [18, 20]. BOsY implements the Bounded Synthesis approach (see the paragraph on the related works in Sec. 1), while STRIX is based on parity games and is the winner of SYNTCOMP 2018, 2019 and 2020. We set a timeout of 180 seconds. The experiments have been run on a 16-cores machine with a 2696.6 MHz AMD core with 62 GB of RAM.

We remark that a comparison with GR(1) synthesis tools is nontrivial. The majority of the tools for GR(1) only support the realizability of the *strict* implication (see for example [8]), not the standard one (which is our case). Therefore, although the latter can be reduced to the former [1], a non-trivial practical effort is required to write an algorithm for this translation.

We considered benchmarks of two types: (i) artificial, and (ii) derived from the SYNTCOMP [14] benchmarks’set. Regarding the artificial benchmarks, we partitioned them in four categories, each containing 30 benchmarks scalable in their dimension  $N$ , for a total of 120 formulas. The categories are the following ones:

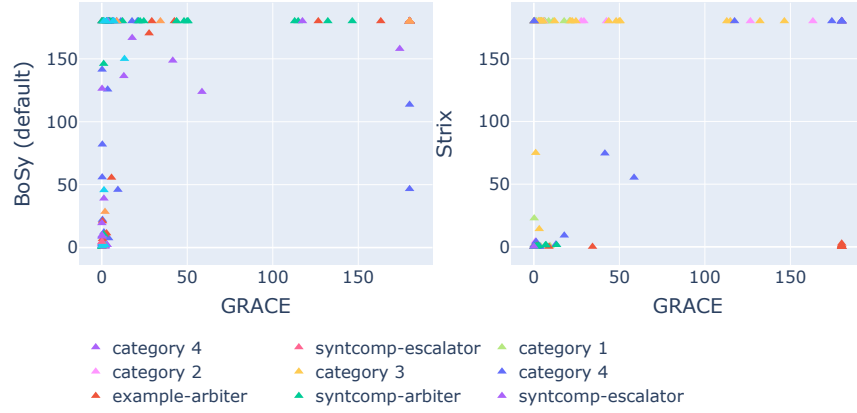
1.  $G(u_0 \rightarrow X(u_1 \rightarrow X(u_2 \rightarrow \dots \rightarrow X(u_N) \dots))) \rightarrow G(\bigwedge_{i=1}^N (u_i \leftrightarrow Xc_i))$
2.  $(G(u_0 \rightarrow X(u_1 \rightarrow X(u_2 \rightarrow \dots \rightarrow X(u_N) \dots))) \wedge X^N G u_N \wedge GF u_N) \rightarrow (\bigwedge_{i=1}^N (u_i \leftrightarrow X^N c_i) \wedge GF c_N)$
3.  $(G(u_0) \wedge XG(u_1) \wedge \dots \wedge X^N G(u_N) \wedge \bigwedge_{i=1}^N GF u_i) \rightarrow (\bigwedge_{i=1}^N G(u_i \leftrightarrow c_i) \wedge \bigwedge_{i=1}^N GF c_i)$
4.  $(\neg u_0 \wedge G^{[0, N]} \neg u_0 \wedge X^{N+1} G u_0) \rightarrow (\bigwedge_{i=1}^N G(u_0 \leftrightarrow Xc_i) \wedge \bigwedge_{i=1}^N GF(c_i \wedge u_0))$

The variables starting with  $u$  are uncontrollable, while those starting with  $c$  are controllables. All the benchmarks are realizable, and were specifically crafted to elicit potential criticalities of GRACE. In particular, the benchmarks in the fourth category have been specifically designed in order to force the minimum  $k$  of the termination of GRACE to increase with their dimension.

Regarding the benchmarks derived from the SYNTCOMP benchmarks’set, we included (i) *simple\_arbiter\_N* (for each  $N \in \{2, 4, 6, 8, 10, 12\}$ ), *escalator\_bidirectional*, which belong to the SYNTCOMP benchmarks’ set, and (ii) our example for an arbiter (Sec. 3.2), with  $N \in \{1, \dots, 15\}$ .


Fig. 2 show the comparison between the tools on all the benchmarks of both types. All times are in seconds. From Fig. 2 (left), it is clear the exponential

<sup>4</sup> <https://es-static.fbk.eu/tools/grace/>



**Fig. 2:** GRACE compared to BOSY (on the left) and to STRIX (on the right)

blowup in solving time in which BOSY occurs. The blowup involves formulas of both types, and of all four categories (of artificial benchmarks). For example, (i) on category 4, the solving times of BOSY on  $N = 13, 14$  are 19.4 and 136.3 sec., respectively, and the corresponding automata have 27 and 31 states, respectively. (ii) on `simple_arbiter_N`, BOSY takes 45.714 sec. for  $N = 8$ , and reaches the timeout for  $N = 10$ . Fig. 3 (left), which compares the dimension of the benchmarks (X axis) with the solving time of GRACE, BOSY and STRIX (Y axis), clearly shows this trend. A more precise study of the complexity of BOSY shows that the majority of the time spent by it is due to the construction of the UCW corresponding to the input formula, which is the task of the tools LTL3BA and SPOT. On the contrary, it is clear from Fig. 2 (left) that GRACE avoids this



**Fig. 3:** On the left, the size of the benchmarks compared to solving time. On the right, GRACE vs BOSY on number of safety sub-problems.

bad behavior, most likely due to the fact that the explicit state `UCW` is never built. Similar considerations can be made for the tool `STRIX` (see Fig. 2, right), except for the category `example-arbiter`, in which the solving times of `STRIX` are consistently better than the ones of `GRACE`. A careful study revealed that all these benchmarks are transformed to the equi-realizable formula `true` by the preprocessor of `OWL` [16] (a tool for  $\omega$ -automata manipulation), which `STRIX` is based on.

The plot in Fig. 3 (right) shows, for each index  $k$  ranging from 1 to 31 (these correspond to the columns), on how many benchmarks (of both types) `GRACE` or `BOSY` terminate with index  $k$  (this corresponds to the height of a column). The benchmarks in category 4 and the ones of `simple_arbiter_N` force `GRACE` to terminate with increasing values of  $k$ . The plot in Fig. 3 points out that `BOSY` does not incur in this growth, except for one benchmark. Nevertheless, the solving times of `GRACE` are still better than the ones of `BOSY`. This witnesses the fact that each safety sub-problem generated by `GRACE` is very simple to solve.

## 7 Conclusions

In this paper, we introduced the logic of `GR-EBR`, an extension of `LTLEBR+P` [4] adding fairness conditions and assumes-guarantees formulas, and studied its realizability problem. We aim at extending the work done in three directions: (i) as far as we know, there are no safety synthesizer (like `SAFETYSYNTH`) that are able to exploit *incrementality*; since in our context, the only part of the automaton that changes between one iteration and the next one is the counter, some work may be saved; (ii) since `GR(1)` is a very efficient fragment, it is important to investigate whether there is a compilation from `GR-EBR` to `GR(1)`; (iii) last but not least, we aim at giving a semantic characterization of `GR-EBR`, and at exploiting the proposed framework for more expressive logics, such as full LTL.

## References

1. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Saar, Y.: Synthesis of reactive (1) designs. *Journal of Computer and System Sciences* **78**(3), 911–938 (2012)
2. Buchi, J.R., Landweber, L.H.: Solving sequential conditions by finite-state strategies. In: *The Collected Works of J. Richard Büchi*, pp. 525–541. Springer (1990)
3. Church, A.: Logic, arithmetic, and automata. In: *Proceedings of the international congress of mathematicians*. vol. 1962, pp. 23–35 (1962)
4. Cimatti, A., Geatti, L., Gigante, N., Montanari, A., Tonetta, S.: Reactive synthesis from extended bounded response LTL specifications. In: *2020 Formal Methods in Computer Aided Design (FMCAD)*. pp. 83–92. IEEE (2020)
5. Cimatti, A., Geatti, L., Gigante, N., Montanari, A., Tonetta, S.: Expressiveness of Extended Bounded Response LTL. *arXiv preprint arXiv:2109.08319* (2021)
6. Claessen, K., Sörensson, N.: A liveness checking algorithm that counts. In: *2012 Formal Methods in Computer-Aided Design (FMCAD)*. pp. 52–59. IEEE (2012)

7. Ehlers, R.: Symbolic bounded synthesis. In: International conference on Computer Aided Verification (CAV). pp. 365–379. Springer (2010)
8. Ehlers, R., Raman, V.: Slugs: Extensible GR(1) synthesis. In: International Conference on Computer Aided Verification. pp. 333–339. Springer (2016)
9. Faymonville, P., Finkbeiner, B., Rabe, M.N., Tentrup, L.: Encodings of bounded synthesis. In: International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). pp. 354–370. Springer (2017)
10. Faymonville, P., Finkbeiner, B., Tentrup, L.: Bosy: An experimentation framework for bounded synthesis. In: International Conference on Computer Aided Verification (CAV). pp. 325–332. Springer (2017)
11. Filiot, E., Jin, N., Raskin, J.F.: An antichain algorithm for LTL realizability. In: International Conference on Computer Aided Verification (CAV). pp. 263–277. Springer (2009)
12. Finkbeiner, B., Hahn, C., Lukert, P., Stenger, M., Tentrup, L.: Synthesizing reactive systems from hyperproperties. In: International Conference on Computer Aided Verification (CAV). pp. 289–306. Springer (2018)
13. Finkbeiner, B., Schewe, S.: Bounded synthesis. *International Journal on Software Tools for Technology Transfer* **15**(5-6), 519–539 (2013)
14. Jacobs, S., Bloem, R.: The 5th reactive synthesis competition (SYNTCOMP 2018)
15. Jacobs, S., Bloem, R., Brenguier, R., Ehlers, R., Hell, T., Könighofer, R., Pérez, G.A., Raskin, J.F., Ryzhyk, L., Sankur, O., et al.: The first reactive synthesis competition (syntcomp 2014). *International journal on software tools for technology transfer* **19**(3), 367–390 (2017)
16. Kretínský, J., Meggendorfer, T., Sickert, S.: Owl: A library for  $\omega$ -words, automata, and LTL. In: Lahiri, S.K., Wang, C. (eds.) *Automated Technology for Verification and Analysis - 16th International Symposium, ATVA 2018, Los Angeles, CA, USA, October 7-10, 2018, Proceedings. Lecture Notes in Computer Science*, vol. 11138, pp. 543–550. Springer (2018). [https://doi.org/10.1007/978-3-030-01090-4\\_34](https://doi.org/10.1007/978-3-030-01090-4_34), [https://doi.org/10.1007/978-3-030-01090-4\\_34](https://doi.org/10.1007/978-3-030-01090-4_34)
17. Kupferman, O., Vardi, M.Y.: Safrless decision procedures. In: 46th Annual Symposium on Foundations of Computer Science (FOCS). pp. 531–540. IEEE (2005)
18. Luttenberger, M., Meyer, P.J., Sickert, S.: Practical synthesis of reactive systems from LTL specifications via parity games. *Acta Informatica* **57**(1), 3–36 (2020)
19. Manna, Z., Pnueli, A.: A hierarchy of temporal properties (invited paper, 1989). In: *Proceedings of the 9th annual ACM symposium on Principles of distributed computing*. pp. 377–410 (1990)
20. Meyer, P.J., Sickert, S., Luttenberger, M.: Strix: Explicit reactive synthesis strikes back! In: *International Conference on Computer Aided Verification*. pp. 578–586. Springer (2018)
21. Piterman, N., Pnueli, A., Sa’ar, Y.: Synthesis of reactive (1) designs. In: *International Workshop on Verification, Model Checking, and Abstract Interpretation*. pp. 364–380. Springer (2006)
22. Pnueli, A., Rosner, R.: On the synthesis of an asynchronous reactive module. In: *International Colloquium on Automata, Languages, and Programming (ICALP)*. pp. 652–671. Springer (1989)
23. Rosner, R.: Modular synthesis of reactive systems. Ph.D. thesis, PhD thesis, Weizmann Institute of Science (1992)
24. Vardi, M.Y., Wolper, P.: Reasoning about infinite computations. *Information and computation* **115**(1), 1–37 (1994)
25. Zhu, S., Tabajara, L.M., Li, J., Pu, G., Vardi, M.Y.: A symbolic approach to safety LTL synthesis. In: *Haifa Verification Conference*. pp. 147–162. Springer (2017)