A Weakness Measure for GR(1) Formulae

Davide G. Cavezza, Dalal Alrajeh, and András György

Imperial College London, United Kingdom
{d.cavezza15,dalal.alrajeh,a.gyorgy}@imperial.ac.uk

Abstract. In spite of the theoretical and algorithmic developments for system synthesis in recent years, little effort has been dedicated to quantifying the quality of the specifications used for synthesis. When dealing with unrealizable specifications, finding the weakest environment assumptions that would ensure realizability is typically a desirable property; in such context the weakness of the assumptions is a major quality parameter. The question of whether one assumption is weaker than another is commonly interpreted using implication or, equivalently, language inclusion. However, this interpretation does not provide any further insight into the weakness of assumptions when implication does not hold. To our knowledge, the only measure that is capable of comparing two formulae in this case is entropy, but even it fails to provide a sufficiently refined notion of weakness in case of GR(1) formulae, a subset of linear temporal logic formulae which is of particular interest in controller synthesis. In this paper we propose a more refined measure of weakness based on the Hausdorff dimension, a concept that captures the notion of size of the omega-language satisfying a linear temporal logic formula. We identify the conditions under which this measure is guaranteed to distinguish between weaker and stronger GR(1) formulae. We evaluate our proposed weakness measure in the context of computing GR(1) assumptions refinements.

1 Introduction

Reactive synthesis is concerned with finding a system implementation that can satisfy a given specification, under all possible environments [35]. When no such implementation exists, a specification is said to be unrealizable [18]. Though there may be many reasons for why a specification is unrealizable, a common cause is an incomplete set of assumptions over the environment behaviour. Several techniques [27,4,15,5] have been proposed in order to compute refinements for incomplete assumptions so as to ensure the realizability of a specification. These approaches consider specifications expressed in a subset of linear temporal logic (LTL), namely generalized reactivity of rank 1 (GR(1)) [12,11,13], for which tractable synthesis methods exist. Their aim is to find the "weakest" assumptions amongst possible alternatives.

Weakness [37] is a feature intended to capture the degree of freedom (or permissiveness) an environment satisfying the assumptions has over its behaviours; generally, weaker assumptions are preferred since they allow for more general

solutions to the synthesis problem [37,17]. Existing approaches formalize the weakness relation between assumptions through logical implication [37,4], i.e., a formula ϕ_1 is weaker than a formula ϕ_2 if $\phi_2 \rightarrow \phi_1$ is valid. However, this notion does not fully capture the weakness concept as permissiveness [14]. Consider the simple example of a bus arbiter whose environment consists of three devices that can request for bus access. Let r_i be the binary signal meaning "device i requests access". An assumption like "device 1 requests access infinitely often" ($\mathbf{GF}r_1$ in LTL) is intuitively less constraining than "device 2 and 3 request access infinitely often" ($\mathbf{GF}(r_2 \wedge r_3)$). However, since the two assumptions refer to disjoint subsets of variables, no implication relation holds between the two.

To enable comparison between such environment assumptions, we propose a quantitative measure for the weakness of GR(1) formulae—based on their interpretation as an ω-language—and a procedure to compute it. The measure builds upon the notion of Hausdorff dimension [39], a quantity providing an indication of the size of an ω-language: the higher the dimension, the wider the collection of distinct ω-words contained in the ω-language. We show that a sufficient condition for assumptions expressed as invariants to be comparable through our measure is the *strong connectedness* of the underlying ω-language. To compare assumptions containing fairness conditions, we identify and measure a language decomposition based on fairness complements. Though we focus on comparing the weakness of assumptions refinements, the scope of its application can be extended to other contexts, e.g., quantitative model checking, in the form of a measure of the set of behaviors violating some given property (see [6] and the Appendix E).

The paper is structured as follows. Related work is presented in Section 2. Notation and background concepts are presented in Section 3. In Section4 we define requirements on a weakness measure in an axiomatic form. In Section 5, we define Hausdorff dimension and explore its relationship with weakness; hence we introduce the proposed weakness measure first for simpler then for generic GR(1) formulae, and provide sufficient conditions guaranteeing its consistency with implication. We also present our refinement of Staiger's algorithm to compute the weakness measure in the GR(1) case. Section 6 presents several applications of our weakness measure to existing GR(1) benchmarks. Omitted details of the experiments and the source code are also available online in [1]. Finally, conclusions are drawn in Section 7. Some proofs are relegated to the appendix.

2 Related Work

The closest notion to our measure is the *entropy* of ω-languages applied by Asarin et al. [6,7] to quantitative model checking. This quantity measures how diverse the ω-words contained in the language of an LTL formula are. However, it is not sufficiently fine-grained to distinguish between weaker and stronger fairness conditions [6]. We will show that our metric based on Hausdorff dimension is capable of making this distinction.

Quality of LTL formulae has also been defined in the context of model verification. The work by Henzinger et al. [25,24] defines a similarity measure between

models of LTL formulae so as to render the model checking output quantitative: instead of returning a true/false response, quantitative model checking computes the distance (*stability radius*) of the model from the boundary of the satisfiability region of an LTL property. The scope of our work is different: the measure we propose can be interpreted as the *extension* of such a satisfiability region, which is independent of a specific model to check against.

An alternative way to measure behaviour sets is via probabilities. Probabilistic model checking [28,23] enhances the syntax and semantics of temporal logics (usually CTL, *computation tree logic*) with probabilities. This allows for the expressions of properties like "the probability of satisfying a temporal logic formula ϕ by the modelled behaviours is at most p." Further extensions of LTL and/or automata with preference metrics alternative to probabilities have been proposed in [9,16,17,3]. The difference between using probabilities/preference metrics and our proposal is that while all of these measures are additional and depend on arbitrary parameters that may not reflect the true weakness of a logical formula, the measure we propose quantifies a concept of weakness *intrinsic* to the LTL formula itself.

The problem of identifying weakest assumptions appears in the context of assume-guarantee reasoning [34,30,19] for compositional model checking. In order to perform model checking of large systems, those systems are generally broken down to components that can be checked independently for correctness. In this context, one of the challenges is to identify the most general (weakest) assumptions over the environment in which each component operates, such that when they are satisfied, the correctness of the entire system is guaranteed. Assumptions are formalized as transition systems (e.g., modal transition systems) rather than declarative LTL specifications, which is the focus of our work.

3 Preliminaries

Languages and Automata. Let Σ be a finite set of symbols, which we call *alphabet*. A *word* over Σ is a finite sequence of symbols in Σ. An *ω-word* is an infinite sequence of such symbols. A set of words is called a *language*, while a set of ω-words is called an ω-language. A word w is explicitly denoted as a sequence of its symbols $w_1 w_2 \dots w_n$, or with a parenthesis notation $(w_1, w_2, \dots, w_n)$, with the symbols separated by commas; the same notation is used for ω-words. The notation w^j denotes the suffix of w starting with w_j.

Given two words v and w, their concatenation is denoted as $v \cdot w$ or simply vw. The same notation is used for the concatenation of a word v and an ω-word w; the concatenation of an ω-word and a word is not defined. Given a set V of finite-length words and a set W of finite-length words or ω-words over the same alphabet Σ, the set $V \cdot W$ is the set of words obtained by concatenating a word in V with a word in W. *Kleene's star operator* yields the set V^* of finite words obtained by concatenating an arbitrary number of words in V. The *omega operator* applied to V yields the set V^ω of ω-words obtained by concatenating a (countably) infinite number of words in V. Naturally, Σ^* and Σ^ω represent,

respectively, the set of all finite words and all ω-words over the alphabet Σ. The star and omega operators can also be applied to single finite-length words, like in w^* and w^ω.

Given an ω-language $L \subseteq \Sigma^\omega$, we denote by $A_n(L)$ the set of all $w \in \Sigma^*$ such that w is a prefix of a word in L and $|w| = n$. We also define $A(L) = \bigcup_{n\in\mathbb{N}} A_n(L)$ the set of all the prefixes of ω-words in L. It is possible to define a topology on Σ^ω. For more details, we refer the reader to [39]. In this context, we only need the notions of closed ω-languages and of their closure. An ω-language L is *closed* if and only if for any ω-word w such that $A(\{w\}) \subseteq A(L)$, $w \in L$. In other words, L is closed if whenever a word w is arbitrarily close (up to a prefix of arbitrary length) to some word in L, then $w \in L$. The *closure* of an ω-language L, denoted by $\mathcal{C}(L)$, is the smallest closed ω-language that contains L.

The notion of regular ω-languages encompasses ω-languages that allow a finite representation through automata. Formally, we define a *regular ω-language* as an ω-language which is accepted by a deterministic Muller automaton. A *deterministic Muller automaton* (DMA) is defined by the quintuple $\mathcal{M} = \langle Q, \Sigma, q_0, \delta, T\rangle$, where Q is a set of states, Σ is the alphabet of the ω-language, q_0 is the initial state, $\delta : Q \times \Sigma \to Q$ is the transition (partial) function and $T \subseteq 2^Q$ is a set (a table) of accepting state sets. Given an ω-word $w \in \Sigma^\omega$, the *run* induced by w onto $\mathcal{M}$ is a sequence of states $\mathcal{M}(w) = q_0 q_1 \dots$ such that q_0 is the initial state and $q_i = \delta(q_{i-1}, w_i)\ \forall i \in \mathbb{N}$. Let $\mathrm{Inf}(w) \subseteq Q$ be the set of states occurring infinitely many times in $\mathcal{M}(w)$. Then an ω-word is said to be *accepted* by $\mathcal{M}$ iff $\mathrm{Inf}(w) \in T$. By extension, the ω-language accepted by $\mathcal{M}$ is the set of ω-words accepted by $\mathcal{M}$.

A *deterministic Büchi automaton* (DBA) $\mathcal{B}$ is defined in the same way as a DMA except for the acceptance condition, which is replaced by a subset of states $F \subseteq Q$. A word w is accepted by $\mathcal{B}$ iff $\mathrm{Inf}(w) \cap F \neq \varnothing$. Given a DBA it is always possible to obtain an equivalent DMA by replacing the Büchi acceptance condition with the table $T = \{Q' \in 2^Q \mid Q' \cap F \neq \varnothing\}$.

Linear Temporal Logic and GR(1). *Linear temporal logic* (LTL) [36] is an extension of Boolean logic with temporal operators. It allows for expressing properties of infinite sequences of assignments to a set $\mathcal{V}$ of Boolean variables. Details of its syntax and semantics are given in Appendix A for completness.

In this paper, we deal with a specific subset of LTL, called *Generalized Reactivity (1)* (GR(1)), which is largely employed in controller synthesis [12,10,27]. This subset makes use of the operators $\mathbf{G}$ ("always"), which states that its operand formula must hold at each step of a valuation sequence, $\mathbf{F}$ ("eventually"), which requires its operand formula to hold at some point in the sequence, and $\mathbf{X}$ ("next"), which states that the operand formula must hold in the state following the one on which the formula is evaluated.

A GR(1) formula over a set of variables $\mathcal{V}$ has the form $\phi = \phi^{\mathcal{E}} \to \phi^{\mathcal{S}}$, where $\phi^{\mathcal{E}}$ and $\phi^{\mathcal{S}}$ are conjunctions of the following units: (i) an *initial condition*, which is a pure Boolean expression over variables in $\mathcal{V}$, denoted by $B^{init}(\mathcal{V})$; (ii) one or more *invariants*, conditions of the form $\mathbf{G}B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V})$, where $B^{\mathrm{inv}}(\mathcal{V} \cup \mathbf{X}\mathcal{V})$ denotes a pure Boolean expression over the set of variables in $\mathcal{V}$ and the set of

atoms obtained by prepending an $\mathbf{X}$ operator to each variable; and (*iii*) one or more *fairness conditions* of the form $\mathbf{GF}B^{fair}(\mathcal{V})$.

The semantics of GR(1), as of LTL, are formalized as ω-*words* over the alphabet $\Sigma = 2^{\mathcal{V}}$. The set of ω-words that satisfy a formula ϕ is a regular ω-language [40] denoted by $L(\phi)$.

4 Problem Statement

In this section, we present an axiomitazation of weakness of an LTL formula. Hereafter, we denote the weakness measure of the LTL formula ϕ as $d(\phi)$: the higher this measure, the weaker ϕ is, i.e., ϕ_2 is weaker than ϕ_1 if $d(\phi_2) \leq d(\phi_1)$.

In settings such as [37,4,2], an LTL formula ϕ_2 is *weaker* than ϕ_1 if and only if $\phi_1 \rightarrow \phi_2$ is valid (that is, it is true for any ω-word). Semantically, this translates to language inclusion: namely, ϕ_2 is weaker than ϕ_1 iff $L(\phi_1) \subseteq L(\phi_2)$. This gives us the first axiom of weakness.

Axiom 1 *Given two LTL formulae ϕ_1 and ϕ_2, if $\phi_1 \rightarrow \phi_2$, then $d(\phi_1) \leq d(\phi_2)$.*

Notice that this criterion defines a partial ordering of specifications: if none of the two formulae implies the other, those are incomparable according to this criterion. However, even for the incomparable case it may be useful to define a preference criterion.

Consider the simple case of two invariants over $\mathcal{V} = \{a, b, c\}$, $\phi_1 = \mathbf{G}(a \wedge b)$ and $\phi_2 = \mathbf{G}c$. Even if the two formulae are incomparable according to implication, i.e., neither one implies the other, it is clear that ϕ_1 allows in some sense fewer behaviors than ϕ_2: at each time step, the former allows for 2 distinct valuations of $\mathcal{V}$ while ϕ_1 allows 4 of them.

Consider the formulae $\phi_3 = \mathbf{G}(a \rightarrow \mathbf{X}b)$ and $\phi_4 = \mathbf{G}((a \wedge b) \rightarrow \mathbf{X}c)$ instead. Despite neither implying the other, we note that ϕ_3 is more restrictive than ϕ_4 asymptotically: that is, for a large enough n, the number of finite prefixes of length n that satisfy ϕ_3 is less than the number of finite prefixes of length n satisfying ϕ_4 ($\#(L(\phi_3)) < \#(L(\phi_4))$). This can be easily understood if one considers that ϕ_3 poses a restriction to the next symbol in an ω-word whenever a is true (which holds in 4 out of 8 possible valuations of $\mathcal{V}$), while ϕ_4 poses a similar restriction when $a \wedge b$ holds (in 2 out of the 8 valuations).

This means that weakness of a formula should be formalized, in addition to Axiom 1, in terms of the number of finite prefixes it allows. Formally:

Axiom 2 *Given two LTL formulae ϕ_1 and ϕ_2, ϕ_1 is said to be weaker than ϕ_2 if there exists some length $\bar{n}$ such that, for every $n > \bar{n}$, the set of prefixes of length n in $L(\phi_1)$ contains more elements than the set of prefixes of the same length in $L(\phi_2)$, i.e., if $\#(A_n(L(\phi_1))) \leq \#(A_n(L(\phi_2)))$, then $d(\phi_1) \leq d(\phi_2)$.*

The last desirable property is that our weakness measure be at least as discriminating as implication in case one formula strictly implies the other.

Axiom 3 *Let ϕ_1 and ϕ_2 be such that $\phi_1 \rightarrow \phi_2$ is valid and $\phi_2 \rightarrow \phi_1$ is not. Then $d(\phi_1) < d(\phi_2)$.*

In the next section, we prove that our proposed weakness measure satisfies Axioms 1 and 2. We then show that, although our weakness measure is not guaranteed to satisfy Axiom 3 in general, we are able to guarantee so for a specific class of formulae.

5 Weakness Measure of GR(1) Formulae

Hausdorff dimension and *Hausdorff measure* are basic concepts in fractal geometry and represent a way to define measures of extension—that is, analogous concepts to length, area, volume from classical geometry—for fractals [22]. Staiger [39] pinpointed a homeomorphism between fractals and regular ω-languages and proposed an analogous interpretation of the two quantities as extension measures of ω-languages. Intuitively, given an ω-language L, its Hausdorff dimension quantifies the growth rate of the number of distinct n-long prefixes of words in the language, over the length n of those prefixes. This makes it a good candidate for quantifying weakness: the less constrained the language is, the more prefixes of a fixed length are contained in it, implying a higher Hausdorff dimension.

The formal definition of Hausdorff dimension is tightly related to the notion of Hausdorff measure. The following definitions are given in [38].

Definition 1 (α-dimensional Hausdorff outer measure). *Given a regular ω-language L over an alphabet Σ with cardinality r, and a nonnegative real value α, the* α-dimensional Hausdorff outer measure *of L is defined as*

$$m_\alpha(L) = \lim_{n\to\infty} \inf_{V\in\mathcal{L}_n} \sum_{v\in V} r^{-\alpha|v|} \tag{1}$$

where $\mathcal{L}_n = \{V \subseteq \Sigma^* \mid V \cdot \Sigma^\omega \supseteq L \text{ and } |v| \geq n \text{ for all } v \in V\}$ is the collection of languages V containing finite words of length at least n and such that every word in L has at least a prefix in V. □

Definition 2 (Hausdorff dimension and measure). *Given an ω-language L, its* Hausdorff dimension*, denoted by* $\dim(L)$*, is the (unique) value $\bar{\alpha}$ such that*

$$m_\alpha(L) = \infty \;\; \alpha < \bar{\alpha}$$
$$m_\alpha(L) = 0 \;\; \alpha > \bar{\alpha}$$

The value $m_{\dim(L)}(L)$ is called the Hausdorff measure *of L.* □

In other words, Hausdorff measure is the limit of a process of approximating the ω-language L by a set V of finite prefixes with length at least n, and weighing each prefix with a quantity $r^{-\alpha|v|}$ that decreases as the prefix length increases. This limit can be finite and positive for at most one value of the α parameter. This value is called Hausdorff dimension.

A related concept appearing in the literature is entropy:

Definition 3 (Entropy [33]). *Given an ω-language $L \subseteq \Sigma^\omega$ over an alphabet of size r, the* entropy *of L is $H(L) = \limsup_{n\to\infty} \frac{1}{n} \log_r \#(A_n(L))$.*

It has been proved [33] that the Hausdorff dimension has a close relationship with the notion of entropy: Specifically, we have $\dim(L) \leq H(L)$ in general, and $\dim(L) = H(L)$ if L is a closed ω-language. Details on how entropy is computed are given in Appendix B.

When L is not closed, the general algorithm presented in [38,39] provides a more refined intuition of what is actually quantified by Hausdorff dimension, which distinguishes it from entropy. The algorithm is based on computing a Muller automaton $\mathcal{M}_L$ accepting L with set of accepting state sets T_L. For each accepting set $S' \in T_L$ and for each state $s \in S'$, consider the ω-language $C_{S'}$ consisting of all the infinite paths in $\mathcal{M}_L$ starting from s and visiting no states outside S'. It can be shown that this language is closed and its entropy $H(C_{S'})$ is independent of the choice of s [38]. The Hausdorff dimension of L is then

$$\dim(L) = \max_{S' \in T_L} H(C_{S'})\,. \tag{2}$$

Hausdorff dimension provides an ordering consistent with the weakness notion defined in Sec. 4. We can interpret it as a measure of the asymptotic degrees of freedom of an ω-language: it quantifies how many different evolutions are allowed to an ω-word once its run remains in an accepting subset of the Muller automaton. The example below shows how it differs from entropy.

Example 1. Consider the LTL formula $\phi_1 = \mathbf{FG}a$ over the variable set $\mathcal{V} = \{a\}$ whose Muller automaton is shown in Fig. 1. The collection of accepting sets to which a state belongs is enclosed in curly braces. Notice that for any $w \in L(\phi_1)$ both valuations of $\mathcal{V}$ are allowed until w reaches the accepting state, and the satisfaction of $\mathbf{G}a$ may be delayed arbitrarily. Therefore, for any finite n, $\#(A_n(L)) = 2^n$, and thereby $H(L(\phi_1)) = 1$.

In this simple DMA, there is only one accepting singleton $\{s_2\}$. Therefore, there is only one $C_{S'} = \{\{a\}^\omega\}$ which allows only the symbol $\{a\} \in 2^{\mathcal{V}}$. This implies $\#(A_n(C_{S'})) = 1$. The Hausdorff dimension is $\dim(L(\phi_1)) = H(C_{S'}) = 0$. This example demonstrates that the Hausdorff dimension isolates the asymptotic behaviour of $L(\phi_1)$ as it depends only on the condition $\mathbf{G}a$ that is eventually satisfied by any ω-word in the ω-language. □

The following theorem shows that Hausdorff dimension is consistent with implication (hence satisfying Axiom 1).

Theorem 1. *Given two LTL formulae ϕ_1 and ϕ_2 such that $\phi_1 \to \phi_2$ is valid,* $\dim(L(\phi_1)) \leq \dim(L(\phi_2))$.

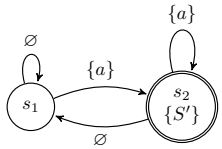

Fig. 1: DMA of $L(\phi_1)$.

Proof. This follows from the language inclusion $L(\phi_1) \subseteq L(\phi_2)$ and the monotonicity of Hausdorff dimension with respect to language inclusion [33].

Note that Theorem 1 does not exclude the situation where one formula strictly implies another, but the two languages have the same Hausdorff dimension, thus violating Axiom 3. We investigate under which conditions this

can or cannot happen in the context of GR(1) formulae and provide a refined weakness measure that bounds the number of cases in which it can happen.

To this end, in what follows, we introduce a new weakness measure for GR(1) based on Hausdorff dimension. We first analyse the dimension of invariants. We then show that under the condition of strong connectedness, it is possible to distinguish between weaker and stronger invariants, in the implication sense (Sec. 5.1). We show how, under the same condition, this measure fails to capture the impact of conjoining a fairness condition (Sec. 5.2). To overcome this, we define a refined weakness measure for GR(1) formulae that comprises two components: the Hausdorff dimension (i) of the whole formula and (ii) of the difference language between the invariant and the fairness conditions (Sec. 5.3).

5.1 Dimension of Invariants

Consider the formula $\phi^{inv} = \mathbf{G}B(\mathcal{V} \cup \mathbf{X}\mathcal{V})$. The ω-language $L(\phi^{inv})$ is closed. Therefore, the Hausdorff dimension of $L(\phi^{inv})$ coincides with its entropy $H(L(\phi^{inv}))$ and can be computed as the maximum eigenvalue of the adjacency matrix of its Büchi automaton (see Appendix B). From this equivalence and Definition 3, it is easy to see that in this case Hausdorff dimension satisfies Axiom 2. In general, Theorem 1 may hold for invariants where one is strictly weaker than the other and both have equal dimensions as demonstrated in the following.

Example 2. Consider the variable set $\mathcal{V} = \{stop\}$ and the formulae $\phi_1^{inv} = \mathbf{G}stop$ and $\phi_2^{inv} = \mathbf{G}(stop \rightarrow \mathbf{X}stop)$. Their Büchi automata are shown in Fig. 2. Clearly $\phi_1^{inv} \rightarrow \phi_2^{inv}$ strictly, however the two languages have the same Hausdorff dimension $\dim\left(L(\phi_1^{inv})\right) = \dim\left(L(\phi_2^{inv})\right) = 0$.

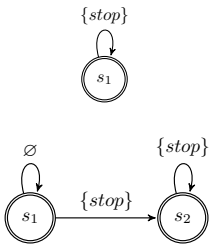

Fig. 2: DBAs of ϕ_1^{inv} (top) and ϕ_2^{inv} (bottom)

There exists, however, a subclass of invariants for which the dimension is strictly monotonic with respect to implication. This subclass is characterized through the concept of strong connectedness of an ω-language. Hereafter, given a word $w \in A(L)$, we denote by $\mathrm{S}_w(L)$ the ω-language formed by the ω-words v such that $wv \in L$ (that is, the suffixes allowed in L after reading w).

Definition 4 (Strongly connected ω-language [33]). *An ω-language L is* strongly connected *if for every prefix $w \in A(L)$ there exists a finite word $v \in \Sigma^*$ such that $\mathrm{S}_{wv}(L) = L$.*

In other words, an ω-language is strongly connected if and only if there exists a strongly connected finite-state automaton which represents it [33], i.e., an automaton such that given any pair of states, each of them is reachable from the other. Using this notion, in the next theorem we provide a sufficient condition over invariants for Axiom 3 to be satisfied (the proof is relegated to Appendix C):

Theorem 2. *Let* $\phi_1^{inv} = \mathbf{G}B_1$ *and* $\phi_2^{inv} = \mathbf{G}B_2$ *be two non-empty invariants such that* $\phi_1^{inv} \rightarrow \phi_2^{inv}$ *is valid,* $\phi_2^{inv} \rightarrow \phi_1^{inv}$ *is not valid and* ϕ_2^{inv} *is strongly connected. Then* $\dim\left(L(\phi_1^{inv})\right) < \dim\left(L(\phi_2^{inv})\right)$.

An interesting kind of invariant that falls in this class is the *one-state invariant*, one that does not use the $\mathbf{X}$ operator: $\phi_s^{inv} = \mathbf{G}B(\mathcal{V})$ whose DBA is shown in Fig. 3. (For succinctness, the set of valuations that label a transition between the same states is denoted by the Boolean expression characterizing it.) In this case, the Hausdorff dimension has a closed form:

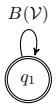

Fig. 3: DBA of a one-state invariant.

$$\dim\left(\phi_s^{inv}\right) = \log_r \#(B(\mathcal{V}))$$

where $r = 2^{\#(\mathcal{V})}$ is the number of valuations of $\mathcal{V}$ and $\#(B(\mathcal{V}))$ is the number of valuations that satisfy $B(\mathcal{V})$. Invariants of this type are clearly strongly connected and satisfy Theorem 2.

Remark 1. Typical examples of GR(1) specifications manually produced, like those of device communication protocols, make use of strongly connected environment assumptions. It is indeed natural to allow environments to be reset to their initial state after some steps. However, when specifications contain "until" operators or response patterns, the procedure to convert them into GR(1) [32] may yield assumptions which are no longer strongly connected. In those cases, a problem similar to that of Example 2 may arise. □

5.2 Fairness and Fairness Complements

Consider the generic fairness condition $\phi^{fair} = \mathbf{GF}B(\mathcal{V})$ whose DBA is shown in Fig. 4. This language is not closed: take a symbol $x \in \Sigma$ that does not satisfy $B(\mathcal{V})$ and the ω-word x^ω. It is clear that $A(\{x^\omega\}) \subseteq A(L(\phi^{fair}))$, but $x^\omega \notin L(\phi^{fair})$. We apply the algorithm in Sec. 5 (cf. equation 2) for non-closed languages. A DMA for $L(\phi^{fair})$ can be obtained from the top DBA in Fig. 4: the accepting sets are $S'_1 = \{q_1, q_2\}$ and $S'_2 = \{q_2\}$. It is easy to see that $H(C_{S'_1}) = 1$ and $H(C_{S'_2}) = \log_r \#(B(\mathcal{V})) \leq 1$. Therefore, $\dim\left(L(\phi^{fair})\right) = 1$, independently of $B(\mathcal{V})$. We conclude that fairness conditions are indistinguishable from the *true* constant, which also has dimension 1.

To allow for a distinction to be made, we characterize the negation of such formula. We call an LTL formula of the kind $\phi^{cfair} = \mathbf{FG}\neg B(\mathcal{V})$ a *fairness complement*. The DMA of $L(\phi^{cfair})$ is shown in the bottom of Fig. 4. The only accepting set is $S' = \{q_2\}$. (Notice that unlike the top one, this automaton accepts only words that stay forever in q_2 from a certain step on.) The language $C_{S'}$ (see Sec. 5) has an entropy of $\log_r \#(\neg B(\mathcal{V}))$. Hence

$$\dim\left(L(\phi^{cfair})\right) = \log_r \#(\neg B(\mathcal{V}))$$

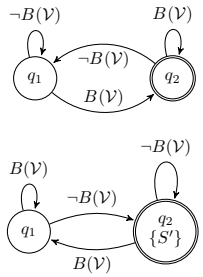

Fig. 4: DBA of $L(\phi^{fair})$ (top) and DMA of $L(\phi^{cfair})$ (bottom).

where $r = 2^{\#(\mathcal{V})}$. Notice that $C_{S'}$ is the language of the formula $\mathbf{G}\neg B(\mathcal{V})$, which is an "asymptotic" condition of ϕ^{cfair}. As observed previously, Hausdorff dimension is strictly monotonic for one-state invariants. Therefore, the weakness of fairness complements can be ranked in terms of the Hausdorff dimension, allowing to compare fairness conditions as follows:

Theorem 3. *Let ϕ_1^{fair} and ϕ_2^{fair} be two fairness conditions such that $\phi_1^{fair} \rightarrow \phi_2^{fair}$ is valid and $\phi_2^{fair} \rightarrow \phi_1^{fair}$ is not valid. Then* $\dim\left(L(\neg\phi_1^{fair})\right) > \dim\left(L(\neg\phi_2^{fair})\right)$.

In other words, the stronger a fairness formula is, the weaker its complement and thereby the higher its dimension.

5.3 Dimension Pairs for GR(1) Formulae

Consider a generic GR(1) formula $\phi = \phi^{init} \wedge \phi^{inv} \wedge \bigwedge_{i=1}^{m} \phi_i^{fair}$. We show through an example that even when ϕ^{inv} is strongly connected, Hausdorff dimension may not distinguish between weaker and stronger fairness conditions in the implication sense. This problem has been also pointed out in [6].

Example 3. Consider the two formulae over the variables $\mathcal{V} = \{a, b\}$: $\phi_1 = \mathbf{G}(a \rightarrow \mathbf{X}b) \wedge \mathbf{GF}a$ and $\phi_2 = \mathbf{G}(a \rightarrow \mathbf{X}b) \wedge \mathbf{GF}b$. The same invariant appears in both, and thereby have the same Hausdorff dimension, but the fairness condition in ϕ_2 is always satisfied when the fairness condition of ϕ_1 is satisfied, by virtue of the invariant itself. However, the ω-word $\{b\}^\omega$ satisfies ϕ_2 but not ϕ_1. So, ϕ_1 implies ϕ_2 but not vice versa.

The language of both formulae is not closed. The Muller automata of ϕ_1 and ϕ_2 are shown at the top and bottom, respectively, in Fig. 5. In both automata, there is an accepting set that covers the entire state space (S'_2 in $L(\phi_1)$ and S'_6 in $L(\phi_2)$. It is possible to show that the maximum $H(C_{S'})$ of equation (2) is achieved exactly for these accepting sets [33,8]. The ω-languages $C_{S'_2}$ in $L(\phi_1)$ and $C_{S'_6}$ in $L(\phi_2)$ both coincide with the language of the invariant alone. Therefore,

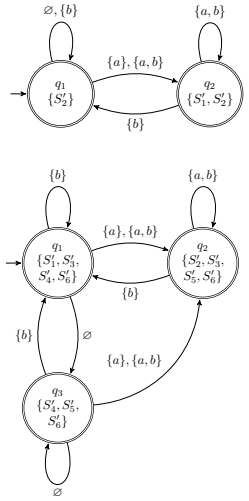

Fig. 5: DMAs of ϕ_1 (top) and ϕ_2 (bottom) of Example 3.

$$\dim(\phi_1) = \dim(\phi_2) = \dim(L(\mathbf{G}(a \rightarrow \mathbf{X}b))) \ .\square$$

To distinguish between the two formulae, we exploit the fact that the complement of a fairness condition is a formula of the kind $\mathbf{FG}B(\mathcal{V})$ which can be compared through Hausdorff dimension. Therefore, we propose a weakness measure which consists of two components: one related to the whole formula and one measuring the ω-language excluded from the invariant by the addition of the fairness conditions.

Definition 5 (Weakness). *The* weakness *of a GR(1) formula* $\phi = (\phi^{init} \wedge \phi^{inv} \bigwedge_{i=1}^{m} \phi_i^{fair})$ *is a pair of numbers* $d(\phi) = (d_1(\phi), d_2(\phi))$ *such that* $d_1(\phi)$ *is the Hausdorff dimension of* $L(\phi)$*; and* $d_2(\phi)$ *is the Hausdorff dimension of* $L(\phi^c) = L(\phi^{init} \wedge \phi^{inv} \wedge \bigvee_{i=1}^{m} \phi_i^{cfair})$*, where* $\phi_i^{cfair} = \neg\phi_i^{fair}$*. The following partial ordering is defined based on the weakness measure: If* $d^i = (d_1^i, d_2^i)$*, with* $i \in 1, 2$ *are weakness measures for two GR(1) formulae, then* $d^1 < d^2$ *if* $d_1^1 < d_1^2$ *or* $d_1^1 = d_1^2$ *and* $d_2^1 > d_2^2$*.*

We apply below this weakness measure to the formulae in Example 3.

Example 4. To compute d_2, let us define $\phi_1^c = \mathbf{G}(a \rightarrow \mathbf{X}b) \wedge \mathbf{FG}\neg a$ and $\phi_2^c = \mathbf{G}(a \rightarrow \mathbf{X}b) \wedge \mathbf{FG}\neg b$. The DMAs of the resulting languages are shown respectively in Fig. 6. Each of them has just one accepting singleton, so the computation of the Hausdorff dimension is straightforward: $\dim(\phi_1^c) = \frac{1}{2}$ and $\dim(\phi_2^c) = 0$. In summary, since ϕ_1 is more restrictive than ϕ_2, the Hausdorff dimension of the ω-language cut out by $\mathbf{GF}a$ is higher than the Hausdorff dimension of the behaviours excluded by $\mathbf{GF}b$. □

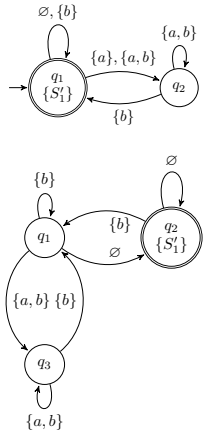

Fig. 6: DMAs of ϕ_1^c (top) and ϕ_2^c (bottom) of Example 4.

The following theorem justifies the use of this dimension pair for weakness quantification when the formulae have the same invariant.

Theorem 4. *Let* $\phi_1 = \phi^{inv} \wedge \bigwedge_{i=1}^{m} \phi_{1,i}^{fair}$ *and* $\phi_2 = \phi^{inv} \wedge \bigwedge_{j=1}^{l} \phi_{2,j}^{fair}$*, such that* $\phi_1 \rightarrow \phi_2$*. Then* $d_1(\phi_1) = d_1(\phi_2)$ *and* $d_2(\phi_1) \geq d_2(\phi_2)$*.*

Proof. Since $\phi_1 \rightarrow \phi_2$, $L(\phi_1) \subseteq L(\phi_2)$. Furthermore, for $i = 1, 2$, $L(\phi_i) = L(\phi^{inv}) \cap L(\bigwedge_{j=1}^{m} \phi_{i,j}^{fair})$. Therefore, $L(\phi^{inv}) \backslash L(\bigwedge_{j=1}^{m} \phi_{1,j}^{fair}) \supseteq L(\phi^{inv}) \backslash L(\bigwedge_{j=1}^{l} \phi_{2,j}^{fair})$, i.e., $L(\phi_1^c) \supseteq L(\phi_2^c)$. Then, by Theorem 1, $\dim(\phi_2^c) \leq \dim(\phi_1^c)$, finishing the proof. □

Therefore, given two formulae with the same invariant, we deem the formula with lower d_2 weaker.

Regarding formulae with the same d_1 and different invariants, we justify heuristically the same order relation. We first note that the Hausdorff dimension of a countable union of ω-languages, as noted in [39], is

$$\dim\left(\bigcup_i L_i\right) = \sup_i \dim(L_i)\,.$$

This property is known as the *countable stability* of Hausdorff dimension. This implies that for any formula ϕ, if $d_2(\phi) \leq d_1(\phi)$ then

$$\dim\left(L(\phi^{inv})\right) = \dim\left(L(\phi) \cup L(\phi^c)\right) = \dim\left(L(\phi)\right)\,.$$

So, if for two formulae, ϕ_1 and ϕ_2, we have $d_1(\phi_1) = d_1(\phi_2) > d_2(\phi_1) > d_2(\phi_1)$, this can be interpreted as the two invariants having the same dimension and the fairness condition of ϕ_1 removing more behaviours than the fairness condition of ϕ_2. In this sense, ϕ_2 is weaker than ϕ_1. This justifies intuitively our weakness definition and the associated partial ordering. In Sec. 6, we illustrate applications of this order relation for comparing GR(1) assumptions.

The computation of $d_2(\phi)$ for a generic ϕ with m fairness conditions can be reduced to the case of a single fairness condition. Based on the countable stability of Hausdorff dimension, we have

$$d_2(\phi) = \sup_{i=1,\ldots,m} d_2(\phi^{init} \wedge \phi^{inv} \wedge \phi_i^{cfair}) \,.$$

Furthermore, the case of a single fairness condition can be further reduced to computing the Hausdorff dimension of an invariant by the following theorem.

Theorem 5. *Given a formula* $\phi^c = \mathbf{G}B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V}) \wedge \mathbf{FG}\neg B^{fair}(\mathcal{V})$ *we have*

$$\dim\left(L(\phi^c)\right) = \dim\left(L(\mathbf{G}(B^{inv} \wedge \neg B^{fair}))\right) \,.$$

P*roof sketch (full proof is presented in Appendix D).* Since $L(\phi^c)$ is not closed, the Hausdorff dimension must be computed from a DMA. The proof (given in Appendix D) consists in showing that the DMA's accepting subsets correspond to the automaton of an ω-language where both B^{inv} and B^{fair} are satisfied at every step. This property is a generalization of the observation made in Sec. 5.2 about the Hausdorff dimension of fairness complements. □

5.4 Initial Conditions

Consider $\phi^{init} = B(\mathcal{V})$. An expression of this form constrains only the first symbol of the ω-words in $L(\phi^{init})$. For the same reason as ϕ^{fair} in Sec. 5.2, $L(\phi^{init})$ is closed, and therefore its dimension can be computed via its entropy. By applying the definition of entropy, it is easy to see that, similarly to the unconstrained language $L(\mathit{true})$,

$$\dim\left(L(\phi^{init})\right) = 1 \,.$$

Consider now a formula $\phi = \phi^{init} \wedge \phi^{inv}$. A DBA $\mathcal{B}$ for $L(\phi)$ can be computed from a DBA $\mathcal{B}_{inv}$ of $L(\phi^{inv})$ by removing all transitions starting from its initial state whose labels do not satisfy $B(\mathcal{V})$. The resulting automaton may leave out parts of $\mathcal{B}_{inv}$ that are no longer reachable from the initial state. This does not happen if $L(\phi^{inv})$ is strongly connected, as in that case any non-initial state in $\mathcal{B}_{inv}$ is reachable from any other state. In this case

$$\dim\left(\phi\right) = \dim\left(\phi^{inv}\right) \,.$$

This implies that the initial conditions do not affect the Hausdorff dimension and hence cannot be always ordered by our weakness measure. This is acceptable since typically, in applications like assumptions refinement, the focus is in assessing invariants or fairness conditions rather than initial conditions [29].

6 Evaluation

We evaluate here the weakness measure through applications to benchmarks within the assumptions refinement domain, demonstrating its usefulness in distinguishing weakness of different formulae, and discussing the computation time bottlenecks. (In Appendix E, we report on our evaluation within another application domain, namely quantitative model checking.)

To this aim, we implemented the weakness measure computation for GR(1) specifications in Python 2.7 and made it publicly available in [1]. Our implementation makes use of the Spot tool [20] for LTL-to-automata conversion. We integrated the weakness computation algorithm within two state-of-the-art counterstrategy-guided assumptions refinement approaches [15,4] (the implementations are available in [1].) Both approaches are based on an iterative procedure of realizability checking and specification refinement: in each iteration, the procedure checks if the current assumptions are sufficient to ensure realizability. If not, a counterstrategy is generated from which a new refinement (i.e., an initial condition, an invariant or a fairness condition that excludes the counterstrategy) is computed and added to the set of available assumptions.

We conducted experiments on two benchmarks for GR(1) assumptions refinement, namely the specifications of a lift controller and of the AMBA-AHB protocol for device communications in its versions for two, four and eight master devices [12,4,29]. The lift controller example specifies a controller for a lift with three floors: the Boolean variable b_i denote the state of the button on floor i; the Boolean variable f_i is true iff the lift is at floor i. For more details on the initial assumptions $\phi^{\mathcal{E}}$ see [4]. The AMBA-AHB protocol provides signals for requesting access to a bus ($hbusreq_i$), for granting access ($hgrant_i$), for signalling the termination of a communication (*hready*), and for identifying the current owner of the bus (*hmaster*). Other signals are detailed in [12]. To our knowledge, the AMBA08 specification is one of the biggest benchmarks available in this field, with 55 binary variables, 28 initial assumptions and 157 guarantees.

Owing to space limitations, we focus below only on our application to [15], and discuss three cases highlighting features of our weakness measure: (*i*) in the first example, we demonstrate the relationship between weakness and implication; (*ii*) second, we consider cases when two formulae are not comparable by implication but can be ranked with our measure; and (*iii*) we discuss the case of formulae equally constraining the environment, which have equal ranking according to our measure. We refer the reader to [1] for the complete results.

Relation between weakness and implication. Consider the lift controller example. Two refinements computed by the automated approach in [15] are: $\phi_1 = \mathbf{G}((\neg b_1 \wedge \neg b_2 \wedge \neg b_3) \rightarrow \mathbf{X}(b_1 \vee b_2 \vee b_3))$; and $\phi_2 = \mathbf{GF}(b_1 \vee b_2 \vee b_3)$. The first forces one of the buttons to be pressed at least every second step in a behaviour. The second forces one of the buttons to be pressed infinitely often in a behaviour. It clear that ϕ_1 implies ϕ_2. We compare the assumptions obtained by refining the original assumptions with the first one and with the second one: $d(\phi^{\mathcal{E}} \wedge \phi_1) = (0.7746, 0)$ and $d(\phi^{\mathcal{E}} \wedge \phi_2) = (0.7925, 0.5)$. Notice that

$d_1(\phi^{\mathcal{E}} \wedge \phi_1) < d_1(\phi^{\mathcal{E}} \wedge \phi_2)$ and this is consistent with the fact that ϕ_1 is stronger than ϕ_2. Consider now the two fairness refinements: $\phi_2 = \mathbf{GF}(b_1 \vee b_2 \vee b_3)$; and $\phi_3 = \mathbf{GF}b_1$. We have $d(\phi^{\mathcal{E}} \wedge \phi_2) = (0.7925, 0.5)$ and $d(\phi^{\mathcal{E}} \wedge \phi_3) = (0.7925, 0.695)$. Here, d_1 is equal for both formulae and $d_2(\phi^{\mathcal{E}} \wedge \phi_2) < d_2(\phi^{\mathcal{E}} \wedge \phi_3)$; this is consistent with the fact that ϕ_2 is weaker than ϕ_3.

Formulae incomparable via implication. Consider ϕ_3 above and $\phi_4 = \mathbf{GF}(b_2 \vee b_3)$. Neither implies the other. However, it is reasonable to argue that ϕ_4 is less restrictive than ϕ_3: while ϕ_3 constrains exactly one button to be pressed infinitely often, ϕ_4 allows the extra choice of which one (out of two) . This intuition is indeed reflected by our computed weakness metric: $d(\phi^{\mathcal{E}} \wedge \phi_3) = (0.7925, 0.695)$ and $d(\phi^{\mathcal{E}} \wedge \phi_4) = (0.7925, 0.5975)$. This expresses the notion that ϕ_4 removes less behaviours from $\phi^{\mathcal{E}}$ than ϕ_3.

Our weakness measure can help spotting asymmetries between assumptions that are syntactically equal but constrain semantically different variables. Consider an extended version of the lift controller example including the input variable *alarm* and the output variable *stop*: whenever *alarm* is set to high, the lift enters a *stop* state where it does not move from the floor it is at. The specification of this system is given in the Appendix F. Computing the weakness of the two refinements $\phi_5 = \mathbf{G}\neg b_1$ and $\phi_6 = \mathbf{G}\neg alarm$ yields $d(\phi^{\mathcal{E}} \wedge \phi^{\mathcal{S}} \wedge \phi_5) = (0.3694, 0.3207)$ and $d(\phi^{\mathcal{E}} \wedge \phi^{\mathcal{S}} \wedge \phi_6) = (0.3746, 0.3346)$. This is consistent with the intuition that the former assumption excludes a part of the desirable system behaviors (all the ones that allow it to reach floor 1), while the latter excludes only the error traces ending in the *stop* state, being then a weaker restriction on the combined behaviors of the controller and the environment.

The following two assumptions refinements are computed for the AMBA-AHB case study with two masters: $\psi_1 = \mathbf{G}(\neg hbusreq_1 \vee \mathbf{X}(hready \vee \neg hbusreq_1))$; and $\psi_2 = \mathbf{G}((\neg hgrant_1 \wedge hready \wedge hbusreq_1) \rightarrow \mathbf{X}(\neg hready \vee \neg hbusreq_1))$. As in the case of the lift example, neither formula implies the other. The weakness of the resulting assumptions is: $d(\psi^{\mathcal{E}} \wedge \psi_1) = (0.9503, 0.9068)$ and $d(\psi^{\mathcal{E}} \wedge \psi_2) = (0.9607, 0.9172)$. The refinement ψ_2 is weaker than ψ_1. Such insight into their weakness could be used to guide the refinement approach (e.g., [4,15]) in choosing to only refine those assumptions that may lead to weaker specifications, for instance further refining ψ_2 rather than ψ_1.

Consistency between equally constraining formulae. Consider the AMBA-AHB protocol with eight masters and the two alternative refinements: $\theta_1 = \mathbf{GF}(hmaster_0 \vee \neg hbusreq_1)$; and $\theta_2 = \mathbf{GF}(hmaster_1 \vee \neg hbusreq_2)$. Clearly the two alternatives express the same kind of constraint on different masters. Since the two masters do not have priorities over each other, expectedly the two refinements have the same weakness: $d(\theta^{\mathcal{E}} \wedge \theta_1) = d(\theta^{\mathcal{E}} \wedge \theta_2) = (0.9396, 0.9214)$.

Performance. The time taken to compute the weakness measure for each refinement (computed via the approach in [15]) was consistently less than 1 minute for the lift controller, AMBA02, and AMBA04 case studies. The time needed on a representative subset of refinements from the AMBA08 example is shown in Fig. 7 as a function of the number of GR(1) conjuncts in the assumptions.

The subset comprises a path from the root of the refinement tree (initial assumptions) to one of the 128 leaves. We observed that 126 of the 128 leaves showed similar performance as the one reported in figure; two of them, instead, took around 6700s. Notice that over 99% of the time is spent on DMA computation, and the remaining time is employed on eigenvalue computation.

When using implication to check whether a formula ϕ_1 implies another formula ϕ_2, it is necessary to produce two automata, one for $L(\phi_1) \cap L(\neg\phi_2)$ and one for $L(\phi_2) \cap L(\neg\phi_1)$, and then run an emptiness check on each of them. When comparing k formulae, this operation must be repeated for $O(k^2)$ pairs of formulae. On the other hand, for a set of formulae containing at most m fairness conditions, our weakness measure requires $m+1$ DMA computations, yielding $O(mk)$ automata needed for comparing k formulae. In this respect, the advantage of our weakness measure resides in the reduced number of DMA computations with respect to implication. The price to pay is that some pairs of formulae distinguishable via implication may produce the same weakness value, as noted above.

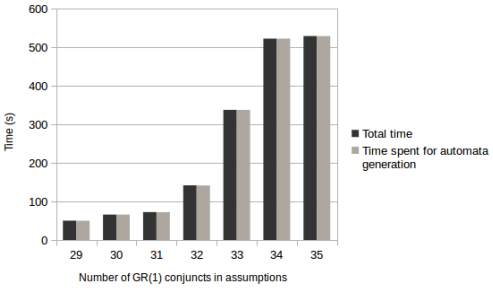

Fig. 7: Execution time of weakness computation for AMBA08

7 Conclusion

In this paper we proposed a new measure for assessing the weakness of GR(1) formulae quantitatively and demonstrated its application in the context of weakest assumptions refinement for GR(1) controller synthesis. We showed that strong connectedness of invariants is a sufficient requirement to guarantee that our measure distinguishes between stronger and weaker formulae in the implication sense. We introduced a component to the measure which allows one to compare formulae with the same dimension based on the weakness of their fairness conditions. The major limitation of the approach is the need for deterministic automata to be produced, which induces high computation time because of the determinization process [21].

As part of our future work, we plan to explore the possibility of refining the weakness relation by including Hausdorff measure in the definition, since Hausdorff measure can distinguish between stronger and weaker ω-languages in case they are not strongly connected [33]. We also intend to investigate algorithms for computing—or approximating at a controlled accuracy—Hausdorff dimension on nondeterministic automata.

Acknowledgments The support of the EPSRC HiPEDS CDT (EP/L016796/1) is gratefully acknowledged.

References

1. https://gitlab.doc.ic.ac.uk/dgc14/WeakestAssumptions
2. Albarghouthi, A., Dillig, I., Gurfinkel, A.: Maximal specification synthesis. ACM SIGPLAN Notices 51(1), 789–801 (2016)
3. Almagor, S., Avni, G., Kupferman, O.: Automatic Generation of Quality Specifications. In: Computer Aided Verification. pp. 479–494 (2013)
4. 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)
5. Alur, R., Moarref, S., Topcu, U.: Pattern-Based Refinement of Assume-Guarantee Specifications in Reactive Synthesis. In: Tools and Algorithms for the Construction and Analysis of Systems. pp. 501–516 (2015)
6. Asarin, E., Blockelet, M., Degorre, A.: Entropy model checking. In: 12th Workshop on Quantitative Aspects of Programming Languages - Joint with European Joint Conference On Theory and Practice of Software (2014)
7. Asarin, E., Blockelet, M., Degorre, A., Dima, C., Mu, C.: Asymptotic behaviour in temporal logic. In: Joint Meeting CSL/LICS. pp. 1–9. ACM Press (2014)
8. Berman, A., Plemmons, R.: Nonnegative Matrices in the Mathematical Sciences. Society for Industrial and Applied Mathematics (1994)
9. Bloem, R., Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Better Quality in Synthesis through Quantitative Objectives. In: Computer Aided Verification, pp. 140–156 (2009)
10. Bloem, R., Cimatti, A., Greimel, K., Hofferek, G., Könighofer, R., Roveri, M., Schuppan, V., Seeber, R.: RATSY A New Requirements Analysis Tool with Synthesis. In: Computer Aided Verification, pp. 425–429 (2010)
11. Bloem, R., Galler, S., Jobstmann, B., Piterman, N., Pnueli, A., Weiglhofer, M.: Specify, Compile, Run: Hardware from PSL. Electronic Notes in Theoretical Computer Science 190(4), 3–16 (2007)
12. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Sa'Ar, Y.: Synthesis of Reactive(1) designs. Journal of Computer and System Sciences 78(3), 911–938 (2012)
13. Braberman, V., D'Ippolito, N., Piterman, N., Sykes, D., Uchitel, S.: Controller synthesis: From modelling to enactment. In: International Conference on Software Engineering. pp. 1347–1350. IEEE (2013)
14. Cassandras, C.G., Lafortune, S.: Introduction to Discrete Event Systems. Springer (2008)
15. Cavezza, D.G., Alrajeh, D.: Interpolation-Based GR(1) Assumptions Refinement. In: Tools and Algorithms for the Construction and Analysis of Systems. pp. 281–297 (2017)
16. Chatterjee, K., De Alfaro, L., Faella, M., Henzinger, T.A., Majumdar, R., Stoelinga, M.: Compositional quantitative reasoning. In: International Conference on the Quantitative Evaluation of Systems. pp. 179–188 (2006)
17. Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Environment Assumptions for Synthesis. In: International Conference on Concurrency Theory. pp. 147–161 (2008)
18. Cimatti, A., Roveri, M., Schuppan, V., Tchaltsev, A.: Diagnostic Information for Realizability. In: International Conference on Verification, Model Checking, and Abstract Interpretation. pp. 52–67 (2008)
19. Cobleigh, J.M., Giannakopoulou, D., Păsăreanu, C.S.: Learning Assumptions for Compositional Verification. In: Tools and Algorithms for the Construction and Analysis of Systems. pp. 331–346 (2003)

20. Duret-Lutz, A., Lewkowicz, A., Fauchille, A., Michaud, T., Renault, E., Xu, L.: Spot 2.0 — a framework for LTL and ω-automata manipulation. In: Automated Technology for Verification and Analysis. vol. 9938, pp. 122–129. Springer (2016)
21. Esparza, J., Ketínský, J., Sickert, S.: From LTL to deterministic automata: A safraless compositional approach. In: Formal Methods in System Design. vol. 49, pp. 219–271 (2016)
22. Falconer, K.: Fractal geometry: mathematical foundations and applications. John Wiley & Sons (2004)
23. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Aspects of Computing 6(5), 512–535 (1994)
24. Henzinger, T.: From Boolean to quantitative notions of correctness. ACM SIGPLAN Notices 45(1), 157 (2010)
25. Henzinger, T.A., Otop, J.: From Model Checking to Model Measuring. In: CONCUR, pp. 273–287 (2013)
26. Horn, R.A., Johnson, C.R. (eds.): Matrix Analysis. Cambridge University Press, New York, NY, USA (1986)
27. Konighofer, R., Hofferek, G., Bloem, R.: Debugging formal specifications using simple counterstrategies. In: Formal Methods in Computer-Aided Design. pp. 152–159 (2009)
28. Kwiatkowska, M.: Quantitative verification: Models, Techniques and Tools. In: Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2015. p. 449. ACM Press (2007)
29. 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)
30. Lomuscio, A., Strulo, B., Walker, N., Wu, P.: Assume-guarantee reasoning with local specifications. International Conference on Formal Engineering Methods pp. 204–219 (2010)
31. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems. Springer (1992)
32. Maoz, S., Ringert, J.O.: GR(1) synthesis for LTL specification patterns. In: Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2015. pp. 96–106. No. 1, ACM Press (2015)
33. Merzenich, W., Staiger, L.: Fractals, dimension, and formal languages. Informatique théorique et applications 28(3-4), 361–386 (1994)
34. Nam, W., Alur, R.: Learning-based symbolic assume-guarantee reasoning with automatic decomposition. Automated Technology for Verification and Analysis pp. 170–185 (2006)
35. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: Principles of Programming Languages. pp. 179–190 (1989)
36. Pnueli, A.: The temporal logic of programs. In: Annual Symposium on Foundations of Computer Science. pp. 46–57 (1977)
37. Seshia, S.A.: Combining Induction, Deduction, and Structure for Verification and Synthesis. IEEE 103(11), 2036–2051 (2015)
38. Staiger, L.: The Hausdorff Measure of Regular ω-languages is Computable. Tech. Rep. August, Martin-Luther-Universität (1998)
39. Staiger, L.: On the Hausdorff measure of regular omega-languages in Cantor space. Tech. Rep. 1, Martin-Luther-Universität Halle-Wittenberg (2015)
40. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. Logics for concurrency pp. 238 – 266 (1996)

Appendix

A LTL Syntax and Semantics

The syntax of LTL is defined by the following grammar:

$$\phi ::= \mathit{true} \mid \mathit{false} \mid p \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \mathbf{F}\phi \mid \mathbf{G}\phi \mid \phi\,\mathbf{U}\phi$$

where $p \in \mathcal{V}$.

The following statements describe LTL semantics, that is when an ω-word is said to satisfy an LTL formula. Hereafter, ϕ and ψ are LTL formulae.

$$\begin{aligned}
w &\models \mathit{true} && \text{always}\\
w &\models \mathit{false} && \text{never}\\
w &\models p && \text{iff } p \in w_1\\
w &\models \neg\phi && \text{iff } w \not\models \phi\\
w &\models \phi \wedge \psi && \text{iff } w \models \phi \text{ and } w \models \psi\\
w &\models \mathbf{X}\phi && \text{iff } w^2 \models \phi\\
w &\models \mathbf{F}\phi && \text{iff } \exists j \in \mathbb{N} \text{ s. t. } w^j \models \phi\\
w &\models \mathbf{G}\phi && \text{iff } \forall j \in \mathbb{N}\ w^j \models \phi\\
w &\models \phi\mathbf{U}\psi && \text{iff } \exists j \in \mathbb{N} \text{ s. t. } w^j \models \psi \text{ and } \forall i < j\ w^i \models \phi
\end{aligned}$$

In other words, **F** can be read as "eventually", **G** as "always", **X** as "next" and **U** as "until".

B Entropy, Automata and Adjacency Matrices

The entropy of a closed language can be computed on its Büchi automaton $\mathcal{B}$ if all states are accepting [33]. This is interpreted as a labelled graph $\mathcal{G} = (Q, E)$, where the set of *nodes* Q is the set of states in $\mathcal{B}$, $E \subseteq Q \times \mathbb{N} \times Q$ is a set of *edges* such that $(q_i, n, q_j) \in E$ if and only if there exist exactly n symbols $\sigma \in \Sigma$ such that $\delta(q_i, \sigma) = q_j$.

Given a subset $Q' \subseteq Q$, the *subgraph* induced by Q' on $\mathcal{G}$ is the graph $\mathcal{G}' = (Q', E')$ such that $(q_i, n, q_j) \in E'$ iff $(q_i, n, q_j) \in E$ and $q_i, q_j \in Q'$.

A graph is *strongly connected* if for any two states $q_i, q_j \in Q$ there exists a path (a sequence of consecutive edges) from q_i to q_j and vice versa. If a graph $\mathcal{G}$ is not strongly connected, it can admit one or more *strongly connected components* (SCCs), which are maximal strongly connected subgraphs of $\mathcal{G}$

A graph can be represented through its *adjacency matrix* A, defined as the square matrix of size $\#(Q)$ whose element in position (i, j) is $A_{ij} = n$ iff $(q_i, n, q_j) \in E$ and $A_{ij} = 0$ if there is no edge connecting q_i and q_j.

The algorithm in [33] to compute entropy is as follows. Given a Büchi automaton with all states accepting and its interpretation as a graph $\mathcal{G}$

1. Compute all SCCs and their adjacency matrices A_i;
2. For every A_i compute the maximum eigenvalue $\rho(A_i)$ (also called *spectral radius* of A_i);
3. Return the maximum $\rho(A_i)$.

C Proof of Theorem 2

Invariants are special cases of safety formulae according to the classification in [31], and therefore can be represented by Büchi automata where every state is accepting. First, we will construct an automaton accepting $L(\phi_2^{inv})$ where each state keeps memory of the last symbol read. The outgoing transitions from each state are labelled only with the valuations that satisfy the invariant. Then we will show that a Büchi automaton of the same form can be obtained for $L(\phi_1^{inv})$ by removing transitions and possibly states from the automaton of $L(\phi_2^{inv})$. This yields an adjacency matrix for $L(\phi_1^{inv})$ with a strictly lower spectral radius (maximum eigenvalue), which corresponds to Hausdorff dimension in our context.

As promised, first we construct the Büchi automata for ϕ_1^{inv} and ϕ_2^{inv}. The construction of $\mathcal{B}$ for $L(\phi^{inv})$ uses a set of states Q which are labelled by a one-to-one function $\lambda : Q \backslash \{q_0\} \to \Sigma$. The transition function is built such that:

1. $\delta(q, \sigma)$ is defined if and only if every ω-word in $\lambda(q)\sigma \cdot \Sigma^\omega$ satisfies $B_2(\mathcal{V} \cup \mathbf{X}\mathcal{V})$;
2. in case $\delta(q, \sigma)$ is defined, $\lambda(\delta(q, \sigma)) = \sigma$.

The initial state q_0 satisfies the following property:

1. for every $\sigma \in \Sigma$, there exists $\delta(q_0, \sigma)$ if and only if there exists $\tau \in \Sigma$ such that $\sigma\tau \cdot \Sigma^\omega \models B$;
2. for every $q \in Q\{q_0\}$, if $q = \delta(q_0, \sigma)$ then $\lambda(q) = \sigma$.

An example is pictured in Figure 8.

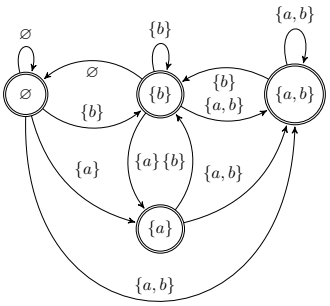

Fig. 8: Büchi automaton of $\mathbf{G}(a \to \mathbf{X}b)$. The state labels $\lambda(q)$ are shown inside the nodes. The initial state (not shown) has transitions towards all the states in the figure, since the first symbol is unconstrained.

We show that an ω-word w satisfies ϕ^{inv} iff it corresponds to an infinite path on $\mathcal{B}$. Suppose $w \models \phi^{inv}$. Then for every $i \in \mathbb{N}$, $w^i \models B(\mathcal{V} \cup \mathbf{X}\mathcal{V})$. Since this formula constrains the first two symbols of w only, any ω-word in $w_i w_{i+1} \cdot \Sigma^\omega$ satisfies B_2. By construction, the automaton $\mathcal{B}$ has a transition $\delta(q_i, w_{i+1}) = q_{i+1}$ such that $\lambda(q_i) = w_i$ and $\lambda(q_{i+1}) = w_{i+1}$. Therefore, if there exists a path from q_0 to q_i induced by the prefix $w_1 \dots w_i$, there exists a path from q_0 to q_{i+1}. As a base case of the induction, consider that $w_1 w_2 \cdot \Sigma^\omega \models B$, and therefore there exists a path from q_0 to q_1 in $\mathcal{B}$.

Conversely, suppose that w is an infinite sequence of transition labels such that $\mathcal{B}(w) = q_0 q_1 \dots$. Then $\delta(q_i, w_{i+1})$ exists for every $i \in \mathbb{N} \cup \{0\}$ and $\lambda(q_i) = w_i$ for all $i \in \mathbb{N}$. By the construction of $\mathcal{B}$ this means that for every $i \in \mathbb{N}$ $w_i w_{i+1} \cdot \Sigma^\omega \models \mathcal{B}$, that implies $w^i \models \phi^{inv}$. Then we can conclude $w \models \phi^{inv}$. This allows us to say that $\mathcal{B}$ is a Büchi automaton of the formula ϕ^{inv}.

If $L(\phi)$ is strongly connected, the $\mathcal{B}$ is also strongly connected, except for the initial state. Let $q_n = \delta(q_0, w) = \delta(\delta(\delta(\dots \delta(q_0, w_1), \dots), w_{n-1}), w_n)$ the state reached by $\mathcal{B}$ after reading the prefix $w \in A_n(L(\phi))$, and $q_m = \delta(q_0, v)$ for $v \in A_m(L(\phi))$. By construction, $\lambda(q_n) = w_n$ and $\lambda(q_m) = v_m$. Since $L(\phi)$ is strongly connected, there exist $v' \in \Sigma^*$ such that $S_{vv'} = L(\phi)$. Therefore $vv'w \in A(L(\phi))$ and $\lambda(q_0, vv'w) = w_n$, so $\delta(q_0, vv'w) = q_n$. So, there exists a path in $\mathcal{B}$ from q_m to q_n. Symmetrically, there exists a path from q_n to q_m. We have proved that for any pair of states reachable from q_0 there is a path between them in both directions, that is the graph induced by non-initial states is strongly connected.

Now consider the automata $\mathcal{B}_1$ and $\mathcal{B}_2$ of $L(\phi_1^{inv})$ and $L(\phi_2^{inv})$ respectively. Since the two formulae are pure invariants, $\phi_1^{inv} \rightarrow \phi_2^{inv}$ if and only if $B_1(\mathcal{V} \cup \mathbf{X}\mathcal{V}) \rightarrow B_2(\mathcal{V} \cup \mathbf{X}\mathcal{V})$. So, given the hypothesis that ϕ_1^{inv} is strictly stronger than ϕ_2^{inv} there must exist a pair of valuations σ, τ such that $\sigma\tau \cdot \Sigma^\omega \models \phi_2^{inv}$ but $\sigma\tau \cdot \Sigma^\omega \not\models \phi_1^{inv}$. By construction this corresponds to at least one transition $\delta(q, \tau)$ that exists in $\mathcal{B}_2$ and does not in $\mathcal{B}_1$. Consequently, we can conclude that $\mathcal{B}_1$ is a proper subgraph of $\mathcal{B}_2$.

The next step is to construct the adjacency matrices corresponding to $\mathcal{B}_1$ and $\mathcal{B}_2$ excluding the respective initial states. Let $Q_1 \backslash \{q_{0,1}\}$ and $Q_2 \backslash \{q_{0,2}\}$ be the set of non-initial states of $\mathcal{B}_1$ and $\mathcal{B}_2$, respectively, and δ_1 and δ_2 their respective transition functions. Let A and B be the adjacency matrices of the graphs $(Q_1 \backslash \{q_0\}, \delta_1)$ and $(Q_2 \backslash \{q_0\}, \delta_2)$. Consider that by construction all transitions between two states are labelled by exactly one valuation: so, each element of these two matrices is either a 0 or a 1.

Since the transitions of δ_2 are a proper subset of the transitions of δ_1, we have for each element (i, j) $A_{ij} \leq B_{ij}$, and $A_{kh} < B_{kh}$ for some (k, h), that is to say $A_{kh} = 0, B_{kh} = 1$. Since ϕ_2^{inv} is strongly connected, its adjacency matrix B is irreducible (for more details on the correspondence between strongly connected digraphs and irreducible matrices see Chapter 6 of [26]). Moreover, $A + B$ is also irreducible, since it is the sum of two nonnegative matrices one of which is irreducible. We can therefore apply the property stated in Chapter 2, Corollary 1.5 of [8], which guarantees that under the given conditions $\rho(A) < \rho(B)$.

Taking the logarithm on both sides, we get $\dim\left(\phi_1^{inv}\right) < \dim\left(\phi_2^{inv}\right)$, finishing the proof. □

D Proof of Theorem 5

The language of $\phi^c = \mathbf{G}B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V}) \wedge \mathbf{FG}\neg B^{fair}(\mathcal{V})$ is not closed. We will therefore build a Muller automaton $\mathcal{M}$ for this language and show that applying the algorithm of Section 4 is equivalent to computing the Hausdorff dimension of the ω-language $L(\mathbf{G}(B^{inv}(\mathcal{V}\cup\mathbf{X}\mathcal{V})\wedge\neg B^{fair}(\mathcal{V})))$. We are supposing that $B^{inv}(\mathcal{V}\cup\mathbf{X}\mathcal{V})$ and $\neg B^{fair}(\mathcal{V})$ are consistent.

Let us first construct a Büchi automaton for $\phi^{inv} = \mathbf{G}B^{inv}(\mathcal{V})$ as in Appendix C. We have shown that any infinite path on this automaton satisfies ϕ^{inv}. We now replace the Büchi winning condition F with a Muller condition T that accounts for satisfying $\phi^{cfair} = \mathbf{FG}\neg B^{fair}(\mathcal{V})$. Let us denote by Q_{cfair} the set of states q such that $\lambda(q)$ satisfies $\neg B^{fair}()$. The accepting table T is then $T := 2^{Q_{cfair}}$.

It is clear that an ω-word w satisfies ϕ^c if and only if w is accepted by $\mathcal{M}$. Suppose w is accepted by $\mathcal{M}$. Then for some $S' \in T$, $\mathit{Inf}(w) = S'$. Since for each $q \in S'$ $\lambda(q)$ satisfies $\neg B^{fair}(\mathcal{V})$, by construction the valuations w_i leading to q satisfies $\neg B^{fair}(\mathcal{V})$. Therefore we conclude that there exists an infinite suffix of w that satisfies $\mathbf{G}\neg B^{fair}(\mathcal{V})$, that is $w \models \phi^{cfair}$. Moreover, w induces an infinite path on the automaton $\mathcal{M}$, and therefore by construction $w \models \phi^{inv}$. Therefore, $w \models \phi$.

Conversely, suppose w satisfies ϕ. Then it satisfies ϕ^{inv}, and thereby induces an infinite path over $\mathcal{M}$. Moreover, it satisfies ϕ^{cfair}, and therefore there exists a suffix of w that satisfies $\mathbf{G}\neg B^{fair}(\mathcal{V})$. So, by construction $\mathit{Inf}(w) \subseteq Q_{cfair}$, that is $\mathit{Inf}(w) \in T$.

The algorithm in Section 4 requires to compute the Hausdorff dimension of every language $C_{S'}$ for $S' \in T$. The Hausdorff dimension of $L(\phi^c)$ is the maximum of such Hausdorff dimensions. The Büchi automaton of the (closed) language $C_{S'}$ corresponds to the subgraph induced by the states in S' onto $\mathcal{M}$, with any state being accepting [38]. The maximum Hausdorff dimension is attained for $S' = Q_{cfair}$.

By construction, $\lambda(q)$ satisfies $\neg B^{fair}(\mathcal{V})$ for every $q \in Q_{cfair}$, and for every pair of consecutive states (q, q') we have $\lambda(q)\lambda(q') \cdot \Sigma^\omega \models B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V})$. Therefore, the subgraph induced by Q_{cfair} corresponds to the Büchi automaton of the closed language $L(\mathbf{G}(B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V}) \wedge \neg B^{fair}(\mathcal{V})))$.

In conclusion,

$$d_2(\phi) = \dim\left(L(\phi^c)\right) = L(\mathbf{G}(B^{inv}(\mathcal{V} \cup \mathbf{X}\mathcal{V}) \wedge \neg B^{fair}(\mathcal{V}))),$$

finishing the proof. □

E Quantitative Model Checking

In this section we provide an application of computing Hausdorff dimension on fairness complements in a quantitative model checking example from [6].

Applied to the model checking problem, our weakness measure extends the quantitative approach in [6] to fairness properties. Consider the Dining Philosophers problem with three philosophers. Let ϕ^{DP} be the GR(1) formula describing all the philosophers' behaviors that do not reach the deadlock state. The goal is to assign a measure to the subset of these behaviors such that none of the philosophers starve. This condition is expressed by the fairness formula $\phi^{DP}_{fair} = \mathbf{GF}(state1 = EAT) \wedge \mathbf{GF}(state2 = EAT) \wedge \mathbf{GF}(state2 = EAT)$. The assigned measure is meant to characterize the degree of satisfaction of this formula by a model of the Dining Philosophers problem. When using entropy only, $H(L(\phi^{DP})) = H(L(\phi^{DP} \wedge \phi^{DP}_{fair})) = 0.0718$. When using our two-component measure function, the degree of satisfaction of a fairness formula is measured indirectly through the subset of behaviors excluded by the formula itself. The result is $d(\phi^{DP}) = (0.0718, 0)$, $d(\phi^{DP} \wedge \phi^{DP}_{fair}) = (0.0718, 0.0479)$. Therefore, contrary to the work in [6], our measure is able to capture the difference in the restrictiveness of the two formulae.

F Specification of the extended lift example

Assumptions:

1. $\neg b_1 \wedge \neg b_2 \wedge \neg b_3 \wedge \neg alarm$
2. $\mathbf{G}((b_1 \wedge f_1) \rightarrow \mathbf{X}\neg b_1)$
3. $\mathbf{G}((b_2 \wedge f_2) \rightarrow \mathbf{X}\neg b_2)$
4. $\mathbf{G}((b_3 \wedge f_3) \rightarrow \mathbf{X}\neg b_3)$
5. $\mathbf{G}((b_1 \wedge \neg f_1) \rightarrow \mathbf{X}b_1)$
6. $\mathbf{G}((b_2 \wedge \neg f_2) \rightarrow \mathbf{X}b_2)$
7. $\mathbf{G}((b_3 \wedge \neg f_3) \rightarrow \mathbf{X}b_3)$

Guarantees:

1. $f_1 \wedge \neg f_2 \wedge \neg f_3 \wedge \neg stop$
2. $\mathbf{G}(\neg(f_1 \wedge f_2) \wedge \neg(f_2 \wedge f_3) \wedge \neg(f_1 \wedge f_3))$
3. $\mathbf{G}((\neg stop \wedge f_3) \rightarrow \mathbf{X}(f_2 \vee f_3))$
4. $\mathbf{G}((\neg stop \wedge f_1) \rightarrow \mathbf{X}(f_1 \vee f_2))$
5. $\mathbf{G}(((f_1 \wedge \mathbf{X}f_2) \vee (f_2 \wedge \mathbf{X}f_3) \vee (f_3 \wedge \mathbf{X}f_2) \vee (f_2 \wedge \mathbf{X}f_1)) \rightarrow (b_1 \vee b_2 \vee b_3))$
6. $\mathbf{GF}((\neg stop \wedge b_1) \rightarrow f_1)$
7. $\mathbf{GF}((\neg stop \wedge b_2) \rightarrow f_2)$
8. $\mathbf{GF}((\neg stop \wedge b_3) \rightarrow f_3)$
9. $\mathbf{G}(stop \rightarrow ((f_1 \rightarrow \mathbf{X}f_1) \wedge (f_2 \rightarrow \mathbf{X}f_2) \wedge (f_3 \rightarrow \mathbf{X}f_3)))$
10. $\mathbf{G}(alarm \rightarrow \mathbf{X}stop)$
11. $\mathbf{GF}f_1$
12. $\mathbf{GF}f_2$
13. $\mathbf{GF}f_3$