

# Debugging Unrealizable Specifications with Model-Based Diagnosis <sup>\*</sup>

Robert Könighofer, Georg Hofferek, and Roderick Bloem

Institute for Applied Information Processing and Communications (IAIK),  
Graz University of Technology, Austria.

**Abstract.** Creating a formal specification for a reactive system is difficult and mistakes happen frequently. Yet, aids for specification debugging are rare. In this paper, we show how model-based diagnosis can be applied to localize errors in unrealizable specifications of reactive systems. An implementation of the system is not required. Our approach identifies properties and signals that can be responsible for unrealizability. By reduction to unrealizability, it can also be used to debug specifications which forbid desired behavior. We analyze specifications given as one set of properties, as well as specifications consisting of assumptions and guarantees. For GR(1) specifications we describe how realizability and unrealizable cores can be computed quickly, using approximations. This technique is not specific to GR(1), though. Finally, we present experimental results where the error localization precision is almost doubled when compared to the presentation of just unrealizable cores.


## 1 Introduction

A formal specification of a reactive system is ideally written before the implementation. This clarifies ambiguities early. Using property synthesis [11, 14, 19, 20, 23, 24], the process of implementing the specification can even be automated. This yields systems which are correct-by-construction. Formal specifications are also used for the precise communication of design intents and interface assumptions. They are even sold as intellectual properties for verification [10]. In all these scenarios, a corresponding implementation is not available when creating the specification.

Just like any other engineering process, writing a formal specification is an error-prone activity [5, 6, 15, 22]. At the same time, debugging an incorrect specification is complicated, especially if no implementation is available. This issue has received little attention in research, and there is almost no tool support available. Ideally, the specification is sound and complete. A specification is *complete* if no incorrect system (with respect to the informal design intent) conforms to it. A specification is *sound*, if all correct systems are valid implementations of the specification. For some applications the specification has to be complete, but

---

<sup>\*</sup> This work was supported in part by the European Commission through projects COCONUT (FP7-2007-IST-1-217069) and DIAMOND (FP7-2009-IST-4-248613).



**Fig. 1.** Debugging an unrealizable specification with model-based diagnosis (MBD).

in any case, it must be sound. This paper addresses the debugging of unsound specifications without a corresponding implementation.

A special case of unsoundness is unrealizability. A specification is unrealizable if no system (no Mealy machine) can implement it. Note that for reactive systems, realizability is not equal to satisfiability [24]. Satisfiability means that there is one input/output trace fulfilling the specification. Realizability requires the existence of a valid output trace for *every* input trace. Additionally, outputs may depend on past and present inputs only. Our experience with the creation of complete specifications [2, 3, 1] shows that mistakes often lead to unrealizable but satisfiable specifications. Manual analysis to find the bug is time consuming or even intractable. Unlike software, there is no way to execute an unrealizable specification in order to locate the error. Debugging specifications that disallow desired behavior (i.e., unsound specifications), as well as specifications which allow undesired behavior (i.e., incomplete specifications) can be reduced to debugging unrealizability of specifications. The reader is referred to [17] for details. Hence, the approach presented in this paper can be used to debug unsoundness as well.

In this paper, we present a technique for error localization in unrealizable specifications of reactive systems, using model-based diagnosis [25] (MBD). In MBD, a system description and some incorrect behavior that contradicts this description are given. The goal is to identify components that may be responsible for the behavior. Our setting is different. We are only given an unrealizable specification. Our first contribution is to define a formalism to make MBD applicable nevertheless: Instead of diagnosing a conflict between the system description and the behavior, we diagnose the unrealizable specification, which is seen as an inherently conflicting system description. We identify properties and signals that can be weakened in order to resolve the unrealizability. Not only single-fault but also multiple-fault diagnoses can be computed. The technique is very generic and applies to most temporal logics, including LTL. It can even be used in other application domains such as the debugging of unsatisfiable formulas in a SAT solver by substituting the concept of realizability with that of satisfiability.

In our solution, diagnoses are computed using minimal unrealizable cores (parts of the specification which are unrealizable on their own) and realizability checks. This can be time consuming. As our second contribution, we therefore show how to boost performance of realizability checking and unrealizable core computation using approximations. We use Generalized Reactivity(1) specifications, but the optimization technique is again rather generic. We finally present experimental results for this class of specifications.

As illustrated in Fig. 1, MBD yields two outcomes: diagnoses and (minimal) conflicts. Diagnoses indicate possible error locations. In order to fix the error in the best way, the user also has to understand the underlying problem. This information is contained in the conflicts, which are unrealizable cores. Conflicts, and thereby the root causes for unrealizability, can be explained with counter-strategies [17]. However, understanding one conflict is not enough since a fix has to resolve all of them. The same problem exists in [5], where one unrealizable core is presented as diagnostic aid. In contrast, our approach identifies exactly those components that can be modified in a way to resolve *all* inconsistencies in the specification simultaneously. This information prevents the user from modifying parts of the specification which contribute to some but not to all conflicts. This work therefore improves and complements [5, 17].

Other related work includes detection of incompleteness in specifications both regarding a given implementation [4, 13, 15] and with the specification analyzed stand-alone [6, 8, 12]. In contrast, our work is focused on specifications that are not sound. MBD has been successfully applied to various software- and hardware debugging problems [7, 21, 27, 26]. The setting considered in this work deviates from the classical setting in that we are not given a system description and a conflicting observation, but only an inherently conflicting system description.

The remainder of this paper is organized as follows. Section 2 explains existing concepts and establishes some notation. Section 3 introduces our debugging method in a generic way. An efficient implementation for Generalized Reactivity(1) specifications is described in Section 4. Experimental results are presented in Section 5. Section 6 concludes the work.

## 2 Preliminaries

### 2.1 Model-Based Diagnosis

*Model-based diagnosis* [9, 25] (MBD) is a technique for error localization in a system. The explanation in this section follows [25]. Let  $\text{SD}$  be a description of the correct system behavior and let  $\text{OBS}$  be an observation of an erroneous behavior. Both  $\text{SD}$  and  $\text{OBS}$  are sets of logical sentences. The system consists of components  $\text{COMP}$ . A component  $c \in \text{COMP}$  can behave abnormally (denoted  $\text{AB}(c)$ ) or normally (denoted  $\neg\text{AB}(c)$ ).  $\text{SD}$  consists of component descriptions of the form  $\neg\text{AB}(c) \Rightarrow N_c$ , where  $N_c$  defines the normal behavior of component  $c$ , and a description of their interplay. The observation contradicts the system description, i.e., it would not be possible if all components behaved normally. Formally, the set  $\text{SD} \cup \text{OBS} \cup \{\neg\text{AB}(c) \mid c \in \text{COMP}\}$  of logical sentences is inconsistent, i.e., contains a contradiction. The goal of MBD is to identify sets of components that may have caused the erroneous behavior  $\text{OBS}$ . Such sets are called diagnoses. Formally,  $\Delta \subseteq \text{COMP}$  is a *diagnosis* iff it is a minimal set such that

$$\text{SD} \cup \text{OBS} \cup \{\neg\text{AB}(c) \mid c \in \text{COMP} \setminus \Delta\} \quad (1)$$

is consistent. Minimal means that no subset  $\Delta' \subset \Delta$  is a diagnosis. A diagnosis  $\Delta$  with  $|\Delta| = 1$  is called a *single-fault diagnosis*. Diagnoses can be computed with

conflicts. A *conflict* is a set  $C \subseteq \text{COMP}$  so that

$$\text{SD} \cup \text{OBS} \cup \{\neg \text{AB}(c) \mid c \in C\} \quad (2)$$

is inconsistent. That is, a conflict is a set of components that cannot all behave normally. A conflict  $C$  is *minimal* if no  $C' \subset C$  is a conflict. A diagnosis must explain all conflicts, so it must have at least one element in common with every conflict. This relation can be formalized with hitting sets. A *hitting set* for a collection  $\mathcal{K}$  of sets is a set  $H$  for which  $\forall K \in \mathcal{K}. H \cap K \neq \emptyset$  holds. A hitting set  $H$  is *minimal* if no subset  $H' \subset H$  is a hitting set. A set  $\Delta \subseteq \text{COMP}$  is a diagnosis iff  $\Delta$  is a minimal hitting set for the set of all conflicts. Equivalently,  $\Delta \subseteq \text{COMP}$  is a diagnosis iff it is a minimal hitting set for the set of all *minimal* conflicts. Hence, computing diagnoses reduces to computing minimal hitting sets for the collection of minimal conflict sets. An algorithm is presented in [25]. It triggers conflict computations on-the-fly. Diagnoses are generated in order of increasing cardinality. That is, simpler explanations are produced first.

## 2.2 Games and the $\mu$ -Calculus

Games and the  $\mu$ -calculus will be used to define realizability and approximations thereof. A (*finite state, two player*) *game* is a tuple  $\mathcal{G} = (Q, \Sigma, T, q_0, \text{Win})$ , where  $Q$  is a finite set of states,  $\Sigma$  is some finite alphabet,  $T : Q \times \Sigma \rightarrow Q$  is a deterministic and complete transition function,  $q_0 \in Q$  is the initial state, and  $\text{Win} : Q^\omega \rightarrow \{\text{false}, \text{true}\}$  is the winning condition. We assume that  $\Sigma = \mathcal{X} \times \mathcal{Y}$ , where  $\mathcal{X} = 2^X$ ,  $\mathcal{Y} = 2^Y$ , and  $X$  and  $Y$  are two disjoint sets of Boolean variables. The game is played by two players, the *environment* and the *system*. A *play*  $\bar{\pi}$  of  $\mathcal{G}$  is defined as an infinite sequence of states  $\bar{\pi} = q_0 q_1 q_2 \dots \in Q^\omega$  such that  $q_{i+1} = T(q_i, \sigma_i)$ . In each step, the letters  $\sigma_i = (x_i, y_i)$  are chosen by the two players in such a way that the environment first chooses an  $x_i \in \mathcal{X}$ , after which the system chooses some  $y_i \in \mathcal{Y}$ .

A play  $\bar{\pi} = q_0 q_1 q_2 \dots$  is *won* by the system if  $\text{Win}(\bar{\pi})$  holds. A state  $q \in Q$  is called *winning*, if starting from this state, the system can, for all possible inputs which it might be presented, choose outputs in each step, such that the play is won. The set of winning states  $W \subseteq Q$  is called the *winning region* of the system.

The (propositional)  $\mu$ -calculus [18] is an extension of propositional logic with a least fixpoint operator  $\mu$  and a greatest fixpoint operator  $\nu$ . We use this calculus, extended with two preimage operators  $\text{MX}^s$  and  $\text{MX}^e$ , to describe fixpoint computations over sets  $Q' \subseteq Q$  of states in a game  $\mathcal{G}$ . Let  $\text{Var}$  be a set of variables ranging over subsets of  $Q$ . Every variable  $Z \in \text{Var}$  and every set  $Q' \subseteq Q$  of states is a valid  $\mu$ -calculus formula. Also, if  $R$  and  $S$  are  $\mu$ -calculus formulas, then so are  $\neg R$ ,  $R \cup S$ , and  $R \cap S$ , with the expected semantics. Finally, for  $Z \in \text{Var}$ , the  $\mu$ -calculus formulas  $\mu Z . R(Z)$ ,  $\nu Z . R(Z)$ ,  $\text{MX}^s(R)$ , and  $\text{MX}^e(R)$

are  $\mu$ -calculus formulas. The latter are defined as

$$\begin{aligned}\mu Z.R(Z) &= \bigcup_i Z_i, \quad \text{where } Z_0 = \emptyset \text{ and } Z_{i+1} = R(Z_i), \\ \nu Z.R(Z) &= \bigcap_i Z_i, \quad \text{where } Z_0 = Q \text{ and } Z_{i+1} = R(Z_i), \\ \text{MX}^s(R) &= \{q \in Q \mid \forall x \in \mathcal{X}. \exists y \in \mathcal{Y}. T(q, (x, y)) \in R\}, \text{ and} \\ \text{MX}^e(R) &= \{q \in Q \mid \exists x \in \mathcal{X}. \forall y \in \mathcal{Y}. T(q, (x, y)) \in R\}.\end{aligned}$$

The operation  $\text{MX}^s(R)$  gives all states from which the system is able to force the play into a state of  $R$  in one step. Analogously,  $\text{MX}^e(R)$  gives all states from which the environment can enforce a visit to  $R$  in one step.

### 2.3 Specifications for Reactive Systems

A reactive system is a Mealy machine that continuously interacts with its environment via a set  $X$  of inputs and a set  $Y$  of outputs. Without loss of generality we assume that all signals are Boolean. We consider two kinds of specifications. The first one, denoted with  $\varphi_P$ , consists of a set  $P$  of properties. Let  $\bar{\sigma} \models p$  denote that trace  $\bar{\sigma} \in \Sigma^\omega$  fulfills property  $p \in P$ . Then  $\bar{\sigma}$  fulfills  $\varphi_P$ , written  $\bar{\sigma} \models \varphi_P$ , iff for all  $p \in P$ ,  $\bar{\sigma} \models p$  holds. We assume two special properties  $\top$  and  $\perp$ , where  $\bar{\sigma} \models \top$  and  $\bar{\sigma} \not\models \perp$  for every  $\bar{\sigma}$ . The second kind of specifications, denoted  $\varphi_{A,G}$ , is given as a pair  $(A, G)$ , where  $A$  is a set of environment assumptions and  $G$  is a set of system guarantees. If all assumptions are met, the specification requires the system to fulfill all guarantees, i.e.,

$$\bar{\sigma} \models \varphi_{A,G} \quad \text{iff} \quad (\text{for all } a \in A, \bar{\sigma} \models a) \text{ implies } (\text{for all } g \in G, \bar{\sigma} \models g). \quad (3)$$

A Generalized Reactivity(1) [23] specification (GR(1), for short) is of the form  $\varphi_{A,G}$ . It can be transformed [23] into a game  $\mathcal{G}^{\text{GR1}} = (Q, \Sigma, T, q_0, \text{Win})$  with  $m$  special sets  $J_i^e \subseteq Q$  and  $n$  sets  $J_j^s \subseteq Q$  of states. Every  $J_i^e$  is a set of accepting states for the environment. They correspond to environment assumptions. Analogously, the sets  $J_j^s$  correspond to system guarantees. The winning region  $W_{\text{sys}}^{\text{GR1}}$  of the system is characterized [23] by the formula

$$W_{\text{sys}}^{\text{GR1}} = \nu Z. \prod_{j=1}^n \mu Y. \bigcup_{i=1}^m \nu X. J_j^s \cap \text{MX}^s Z \cup \text{MX}^s Y \cup \neg J_i^e \cap \text{MX}^s X. \quad (4)$$

A GR(1) specification is realizable iff  $q_0 \in W_{\text{sys}}^{\text{GR1}}$  in the corresponding game.

### 2.4 Minimization Algorithms

We will use two different minimization algorithms for unrealizable core computation. *Delta Debugging* [28] is an algorithm to isolate the cause of a failure. Given a procedure `test` and some input  $C$  that makes `test` fail (denoted `test(C) = ✕`),

it computes a minimal input  $\hat{C} = \text{ddmin}^{\text{test}}(C)$ , with  $\hat{C} \subseteq C$ , which still makes `test` fail. In the best case, `ddmin` requires only a logarithmic number of checks whether `test` fails. In the worst case, a quadratic number of checks is needed. The reader is referred to [28] for detailed information on the algorithm.

The second minimization algorithm [5] we use removes one element after the other from the given set  $C$ . If `test` passes without a particular element, this element is part of the unrealizable core and is added again. Otherwise, the algorithm proceeds without this element. Thus, exactly  $|C|$  invocations of `test` are required. This algorithm will be denoted  $\text{linMin}^{\text{test}}$ .

### 3 Model-Based Diagnosis for Unrealizability

In this section, we describe how MBD can be used to perform error localization in an unrealizable specification. The difference to standard MBD is that we do not diagnose a conflict between a system description and an observation but an inherently conflicting system description. Our technique can be applied to many specification languages. The prerequisites are that properties can be removed from specifications, output signals can be existentially quantified in properties, and that a decision procedure for the realizability of a specification is available.

#### 3.1 Property Specifications

In this section, we assume that the unrealizable specification is given as a set  $P$  of properties for a system with inputs  $X$  and outputs  $Y$ . Since every property typically represents a relatively self-contained and independent aspect of the system behavior, we define every property to be a component. That is,  $\text{COMP}_P = P$ . Furthermore, we define the system description  $\text{SD}_P$  to be the tuple  $(P, X, Y)$ . There is no observation in our setting. As an abstraction for the notion of logical consistency, we define a function  $\text{consistent}_{\text{SD}_P} : 2^{\text{COMP}_P} \rightarrow \{\text{true}, \text{false}\}$ . Intuitively,  $\text{consistent}_{\text{SD}_P}(P')$  with  $P' \subseteq \text{COMP}_P$  gives `true` iff the system description  $\text{SD}_P$  is consistent under the assumption that all components  $c \in P'$  behave normally. More formally, we transform  $\varphi_P$  to  $\tilde{\varphi}_P$  such that for every trace  $\bar{\sigma} \in \Sigma^\omega$ ,  $\bar{\sigma} \models \tilde{\varphi}_P$  iff for all  $p \in P$ ,  $\neg \text{AB}(p)$  implies  $\bar{\sigma} \models p$ . Moreover,  $\text{consistent}_{\text{SD}_P}(P') = \text{true}$  iff  $\tilde{\varphi}_P$  is consistent with  $\neg \text{AB}(p')$  for all  $p' \in P'$ . That is, abnormal properties do not have to be fulfilled in  $\tilde{\varphi}_P$ , so we can think of them as removed from the specification. Let `realizable` be a function which decides the realizability of a specification. Then, we can finally use the notion of realizability to define consistency as  $\text{consistent}_{\text{SD}_P}(P') = \text{realizable}(P')$ .

**Lemma 1.** *The function  $\text{consistent}_{\text{SD}_P}$  is monotonic, i.e.,  $\forall P'' \subseteq P' \subseteq P. \neg \text{consistent}_{\text{SD}_P}(P'') \Rightarrow \neg \text{consistent}_{\text{SD}_P}(P')$ .*

This follows from the fact that adding properties to an unrealizable specification preserves unrealizability. In analogy to Eq. 2, a conflict is a set  $C_P \subseteq \text{COMP}_P$

such that  $\text{consistent}_{\text{SD}_P}(C_P) = \text{false}$ , i.e.,  $\text{realizable}(C_P) = \text{false}$ . A minimal conflict is therefore what is also called a minimal unrealizable core of the specification [5]. Corresponding to Eq. 1, a diagnosis is finally a minimal set  $\Delta_P \subseteq \text{COMP}_P$  such that  $\text{consistent}_{\text{SD}_P}(\text{COMP}_P \setminus \Delta_P) = \text{true}$ .

**Observation 1** *Let  $C_P$  be a conflict for an unrealizable specification  $\varphi_P = P$ . In order to make  $\varphi_P$  realizable, at least one property  $p \in C_P$  must be modified.*

One cannot obtain a realizable specification by modifying properties  $p \in (P \setminus C_P)$  only, because the set  $C_P$  of properties forms an unrealizable specification already. Any superset of properties will be unrealizable as well (cf. Lemma 1). Informally speaking, there must be something wrong with at least one conflict element.

**Observation 2** *A diagnosis  $\Delta_P \subseteq P$  is a minimal set of properties that can be modified in such a way that the specification becomes realizable.*

Clearly, the specification becomes realizable if all  $p \in \Delta_P$  are replaced by  $\top$ . Minimality follows from the definition.

*Example 1.* The specification  $\varphi_P = \{p_1, p_2, p_3\}$  with<sup>1</sup>  $p_1 = \text{G}(I_1=0 \Rightarrow O_1=1)$ ,  $p_2 = \text{G}(I_1=1 \Rightarrow O_1=0)$ ,  $p_3 = \text{GF}(I_1 \Leftrightarrow O_1 \wedge I_2 \Leftrightarrow O_2)$ ,  $X = \{I_1, I_2\}$ , and  $Y = \{O_1, O_2\}$  is unrealizable because  $p_1$  and  $p_2$  require that output  $O_1$  is always the negation of input  $I_1$ , while  $p_3$  requires these two signals to be equal infinitely often. The specification contains the conflicts  $C_1 = \{p_1, p_2, p_3\}$ ,  $C_2 = \{p_1, p_3\}$ , and  $C_3 = \{p_2, p_3\}$ , the latter two being minimal. Conflict  $C_2$  states that no system can implement  $p_1$  and  $p_3$  at the same time, independent of  $p_2$ . Analogously for  $C_3$ . The sets  $\Delta_1 = \{p_3\}$  and  $\Delta_2 = \{p_1, p_2\}$  are the minimal hitting sets for the collection of all minimal conflicts, since they share at least one element with both  $C_2$  and  $C_3$ . Hence,  $\Delta_1$  and  $\Delta_2$  are (the only) diagnoses. According to Observation 2, the user can therefore weaken either property  $p_3$  or both properties  $p_1$  and  $p_2$  in order to make the specification realizable. Suppose now that the user is only given the unrealizable core  $C_2$ . In this case, she might attempt to modify  $p_1$  alone in order to resolve the unrealizability. However, this is not possible because no matter how  $p_1$  is changed, even if it is  $\top$ , there is still a conflict between  $p_2$  and  $p_3$ . Our debugging approach takes this circumstance into account by combining information about all unrealizable cores. This allows for more precise error localization than the presentation of a single unrealizable core.

### 3.2 Assumption/Guarantee Specifications

Let  $\varphi_{A,G} = (A, G)$  be an unrealizable specification consisting of the environment assumptions  $A$  and the system guarantees  $G$ .

**Proposition 1.** *If a diagnosis is defined to be a minimal set  $\Delta_{A,G} \subseteq (A \cup G)$  of assumptions and guarantees which can be modified in such a way that the specification becomes realizable, then every set  $\{a\}$  for  $a \in A$  is a diagnosis.*

<sup>1</sup> We use the common LTL syntax, where  $\text{G}$  represents the temporal operator “always” and  $\text{F}$  denotes “eventually”.

This follows directly from Eq. 3, since replacing any assumption with  $\perp$  would give a realizable specification. In other words, it does not make sense to search for assumptions that can be modified in order to obtain a realizable specification, because every assumption can be modified in such a way. But not every guarantee can be altered to obtain a realizable specification. Hence, we define  $\text{COMP}_G = G$ , implicitly assuming that all assumptions are correct. Furthermore, we define the system description  $\text{SD}_{A,G}$  to be the tuple  $(A, G, X, Y)$ , and  $\text{consistent}_{(A,G,X,Y)}(G') = \text{realizable}((A, G'))$  for  $G' \subseteq G$ . It follows that a conflict is a set  $C_G \subseteq G$  of guarantees such that  $\text{realizable}((A, C_G)) = \text{false}$ . Consequently, a diagnosis is a minimal set  $\Delta_G \subseteq G$  such that  $\text{realizable}((A, G \setminus \Delta_G)) = \text{true}$ .

Lemma 1 and the Observations 1 and 2 apply in this case as well, with the obvious adaptations. Since specifications  $\varphi_P$  are special cases of specifications  $\varphi_{A,G}$  (with  $A = \emptyset$ ), we restrict our investigations to the latter form in the following.

### 3.3 Diagnosing Variables

In the previous section we showed how to identify (sets of) guarantees that can be weakened in order to obtain a realizable specification. In this section we define a formalism to identify *signals* that may be over-constrained, i.e., signals that can be less restricted in order to resolve the unrealizability. We also show how the two approaches can be combined.

Let  $\varphi_{A,G} = (A, G)$  be an unrealizable specification over the inputs  $X$  and outputs  $Y$ , and let  $\bar{\sigma} = (x_0, y_0)(x_1, y_1) \dots \in (\mathcal{X} \times \mathcal{Y})^\omega$  be a trace. As already introduced rather informally in [17], we can define an existential quantification  $(A, \exists Y'. G)$  of the outputs  $Y' \subseteq Y$  in the guarantees of  $\varphi_{A,G}$  with semantics

$$\bar{\sigma} \models (A, \exists Y'. G) \quad \text{iff} \quad (\forall a \in A. \bar{\sigma} \models a) \text{ implies } (\forall g \in G. \bar{\sigma} \models \exists Y'. g).$$

The existential quantification  $\exists Y'. g$  in one single guarantee  $g \in G$  is defined as

$$\bar{\sigma} \models \exists Y'. g \quad \text{iff} \quad \exists y'_0 y'_1 y'_2 \dots \in \left(2^{Y'}\right)^\omega. (x_0, y_0^E)(x_1, y_1^E)(x_2, y_2^E) \dots \models g,$$

where  $y_i^E = (y_i \setminus Y') \cup y'_i$  for all  $i \geq 0$ . Informally speaking, an existential quantification  $\exists Y'. G$  of the variables  $Y' \subseteq Y$  in all guarantees  $g \in G$  removes all restrictions on the variables  $y \in Y'$ . The specification  $(A, \exists Y'. G')$  allows arbitrary values for all outputs  $y \in Y'$  in all time steps. Also note that the quantification is performed on every single guarantee, and not on the conjunction of all guarantees. With  $\text{COMP}_Y = Y$  and  $Y' \subseteq Y$ , we define  $\text{consistent}_{(A,G,X,Y)}(Y') = \text{realizable}((A, \exists Y \setminus Y'. G))$ . Consequently, a conflict is a set  $C_Y \subseteq Y$  of outputs such that  $\text{realizable}((A, \exists Y \setminus C_Y. G)) = \text{false}$ . Finally, a diagnosis is a minimal set  $\Delta_Y \subseteq Y$  such that  $\text{realizable}((A, \exists \Delta_Y. G)) = \text{true}$ . Every  $\Delta_Y$  contains signals that may be over-constrained, because removing restrictions on these signals resolves the unrealizability.

An alternative definition  $\exists$  of the existential quantification operates on the conjunction of all guarantees instead of every guarantee in isolation. As an example where this makes a difference, consider the specification  $\varphi_{A,G} = (A, G) =$



( $\{\top\}, \{G(O), G(\neg O)\}$ ), where  $O$  is the only output. We have that  $(A, \exists\{O\} . G)$  is realizable (it allows all traces), while  $(A, \exists\{O\} . G)$  is not. Consequently,  $\{O\}$  is a diagnosis when using  $\exists$  but not when using  $\exists$ . Our approach works for both definitions. However, we decided for  $\exists$  because we think that  $\{O\}$  should be a diagnosis for the example. After all, if there was no output  $O$ , there would also be no conflict. Furthermore, the user can often comprehend what a quantification in one guarantee means. Understanding what a quantification in the conjunction of all guarantees means is typically much more difficult, because complex dependencies between the guarantees may exist.

The approaches of diagnosing properties and signals can also be combined by defining  $\text{COMP} = G \cup Y$ . With  $\text{consistent}_{(A,G,X,Y)}(B) = \text{realizable}((A, \exists Y \setminus B . (G \cap B)))$  and  $B \subseteq \text{COMP}$  we have that  $C_{Y,G} \subseteq (Y \cup G)$  is a conflict iff  $\text{realizable}((A, \exists Y \setminus C . (G \cap C))) = \text{false}$ . A diagnosis is now a minimal set  $\Delta_{Y,G} \subseteq (Y \cup G)$  such that  $\text{realizable}((A, \exists(Y \cap \Delta) . (G \setminus \Delta))) = \text{true}$ .

**Theorem 1.** *Every diagnosis  $\Delta_Y$  and every diagnosis  $\Delta_G$  for an unrealizable specification  $\varphi_{A,G}$  is also a diagnosis with respect to the definition of  $\Delta_{Y,G}$ .*

*Proof.*  $\text{realizable}((A, \exists \Delta_Y . G)) \Rightarrow \text{realizable}((A, \exists(Y \cap \Delta_Y) . (G \setminus \Delta_Y)))$  and  $(\forall \Delta'_Y \subset \Delta_Y . \neg \text{realizable}((A, \exists \Delta'_Y . G))) \Rightarrow (\forall \Delta'_Y \subset \Delta_Y . \neg \text{realizable}((A, \exists(Y \cap \Delta'_Y) . (G \setminus \Delta'_Y))))$  since  $Y \cap \Delta'_Y = \Delta'_Y$  and  $G \setminus \Delta'_Y = G$  for all  $\Delta'_Y \subseteq \Delta_Y$ . Analogously for  $\Delta_G$ .

Theorem 1 states that the definition of  $\Delta_{Y,G}$  subsumes those of  $\Delta_Y$  and  $\Delta_G$ . Having all diagnoses  $\Delta_{Y,G}$ , we would not gain further diagnoses by computing  $\Delta_Y$  and  $\Delta_G$ . Hence, we will use the definition of  $\Delta_{Y,G}$  in the following.

*Example 2.* Let  $\varphi_{A,G} = (\emptyset, \{p_1, p_2, p_3\})$ , where  $p_1, p_2$ , and  $p_3$  are defined as for Example 1. The minimal conflicts are  $C_1 = \{p_1, p_3, O_1\}$  and  $C_2 = \{p_2, p_3, O_1\}$ . The sets  $\Delta_1 = \{p_3\}$ ,  $\Delta_2 = \{p_1, p_2\}$ , and  $\Delta_3 = \{O_1\}$  form the minimal hitting sets for the collection  $\{C_1, C_2\}$  of minimal conflicts, and hence also the diagnoses for the unrealizable specification  $\varphi_{A,G}$ . Compared to Example 1, we obtain  $\Delta_3 = \{O_1\}$  as additional diagnosis. That is, when using not only guarantees but also outputs as components for diagnosis, we also get the explanation that  $O_1$  may be over-constrained. No diagnosis involves  $O_2$ , so  $O_2$  does not contribute to the unrealizability. Properties are formulated in terms of signals. Therefore, having both properties and signals as diagnoses allows to track down the error from both these dimensions. For instance, in this example, it is natural to assume that there is something wrong with  $O_1$  in  $p_3$  but not with  $O_2$  in  $p_3$ .

### 3.4 Implementation

An implementation of the diagnosis approach as described in the previous section is straightforward. The only prerequisite is that a decision procedure for the realizability of a specification has to be available. We use the algorithm of [25] to compute diagnoses via a hitting set tree. It requires a procedure to compute conflicts not containing a given list of components, if such a conflict exists. This

can be implemented with a single realizability check. However, the algorithm performs better if the computed conflicts are minimal. Such a procedure can be implemented as

$$\text{cNot}_{\text{SD}}(B) = \begin{cases} \text{None} & \text{if } \text{consistent}_{\text{SD}}(\text{COMP} \setminus B) \\ \text{min}_{\text{SD}}(\text{COMP} \setminus B) & \text{otherwise} \end{cases},$$

where  $\text{min}_{\text{SD}}(M)$ , with  $M \subseteq \text{COMP}$  and  $\neg \text{consistent}_{\text{SD}}(M)$ , returns a set  $\hat{M} \subseteq M$  such that  $\neg \text{consistent}_{\text{SD}}(\hat{M})$  and  $\forall M' \subset \hat{M}. \text{consistent}_{\text{SD}}(M')$ . The procedure  $\text{min}_{\text{SD}}$  can be implemented, e.g., as  $\text{min}_{\text{SD}}(M) = \text{ddmin}^{\text{test}_{\text{SD}}}(M)$  or  $\text{min}_{\text{SD}}(M) = \text{linMin}^{\text{test}_{\text{SD}}}(M)$  with  $\text{test}_{\text{SD}}(M') = \mathbf{x} \Leftrightarrow \neg \text{consistent}_{\text{SD}}(M')$ . Experiments [17] show that  $\text{ddmin}$  is often much faster than  $\text{linMin}$ . When using  $\text{ddmin}$ , the monotonicity of  $\text{test}$  (cf. Lemma 1) can be exploited to speed up computation [28]: all encountered sets  $M'$  for which  $\text{test}_{\text{SD}}(M') \neq \mathbf{x}$  holds are stored. If a subset  $M''$  of a stored set  $M'$  is tested,  $\text{test}_{\text{SD}}(M'') \neq \mathbf{x}$  can be returned without actually invoking the realizability check.

## 4 Efficient Implementation for GR(1) Specifications

In our framework, diagnoses computation requires many unrealizable core computations, which in turn require lots of realizability checks. Thus, it is of utmost importance that these procedures are implemented efficiently. In this section, we show how the performance of these operations can be improved, using GR(1) specifications as an example. For GR(1) specifications, realizability can be decided by constructing a game, computing the winning region for the system, and checking if the initial state is contained in this winning region. This applies to many other specification languages as well. We use approximations of the winning region to define approximations of realizability. The only GR(1)-specific part is how these approximations of the winning region are defined.

A procedure  $\text{realizable}_O(\varphi)$  is called an *over-approximation* of  $\text{realizable}(\varphi)$  iff  $\forall \varphi. \text{realizable}(\varphi) \Rightarrow \text{realizable}_O(\varphi)$ . Such an over-approximation can be used to compute a minimal unrealizable core in two steps. In the first step, an over-approximation  $C'$  of the core is computed using a minimization algorithm which repeatedly applies  $\text{realizable}_O$ . In a second step,  $C'$  is further reduced to the exact core  $C$  by repeatedly applying the exact realizability check. If a procedure  $\text{realizable}_O$  can be found which is both fast and accurate, this two-step approach can increase the performance of the unrealizable core computation significantly. The leverage comes from the fact that the expensive exact checks are performed on relatively small subsets of the specification only.

For GR(1) specifications, we can define  $\text{realizable}^{\text{GR1}}(\varphi) \Leftrightarrow q_0 \in W_{\text{sys}}^{\text{GR1}}$ . See Section 2.3. In order to obtain approximations to this procedure, we define the

following sets of states in the GR(1) game  $\mathcal{G}^{\text{GR1}} = (Q, \Sigma, T, q_0, \text{Win})$ :

$$\begin{aligned}
A_{\text{sys}}^{\text{GR1}} &= \mu X . A' \cup \text{MX}^s(X) \quad \text{with} \quad A' = \bigcup_{i=1}^m \nu Y . \neg J_i^e \cap \text{MX}^s(Y), \\
B_{\text{sys}}^{\text{GR1}} &= \mu X . B' \cup A_{\text{sys}}^{\text{GR1}} \cup \text{MX}^s(X) \quad \text{with} \quad B' = \nu Y . \text{MX}^s(Y) \cap \bigcap_{j=1}^n J_j^s, \text{ and} \\
C_{\text{env}}^{\text{GR1}} &= \mu X . C' \cup \text{MX}^e(X) \quad \text{with} \quad C' = \bigcup_{j=1}^n \nu Y . \text{MX}^e(Y) \cap \neg J_j^s \cap \bigcap_{i=1}^m J_i^e.
\end{aligned}$$

These sets were chosen such that they are close to  $W_{\text{sys}}^{\text{GR1}}$  or  $Q \setminus W_{\text{sys}}^{\text{GR1}}$ , yet fast to compute. Some more approximations were tried, but the listed ones showed the best performance in experiments.  $A_{\text{sys}}^{\text{GR1}}$  contains all states from which the system can force the environment to violate one assumption. The set  $B'$  consists of all states from which the system can enforce to stay in the intersection of all  $J_j^s$  forever, thus fulfilling all system guarantees.  $B_{\text{sys}}^{\text{GR1}}$  contains the set of states from which the system can force the play into a state of  $B'$  or  $A_{\text{sys}}^{\text{GR1}}$ . Clearly,  $A_{\text{sys}}^{\text{GR1}} \subseteq W_{\text{sys}}^{\text{GR1}}$  and  $B_{\text{sys}}^{\text{GR1}} \subseteq W_{\text{sys}}^{\text{GR1}}$ .  $C'$  is the set of states from which the environment can enforce to stay in the intersection of all sets  $J_i^e$  of accepting states of the environment but outside one particular set  $J_j^s$  of accepting states of the system. Hence,  $C_{\text{env}}^{\text{GR1}} \subseteq (Q \setminus W_{\text{sys}}^{\text{GR1}})$ . With these sets, we define  $\text{realizable}_O^{\text{GR1}}(\varphi) \Leftrightarrow q_0 \notin C_{\text{env}}^{\text{GR1}}$  as an over-approximation of  $\text{realizable}^{\text{GR1}}$ , and

$$\text{realizable}_E^{\text{GR1}}(\varphi) = \begin{cases} \text{true} & \text{if } q_0 \in A_{\text{sys}}^{\text{GR1}} \\ \text{true} & \text{else if } q_0 \in B_{\text{sys}}^{\text{GR1}} \\ \text{false} & \text{else if } q_0 \in C_{\text{env}}^{\text{GR1}} \\ \text{realizable}^{\text{GR1}}(\varphi) & \text{otherwise} \end{cases}$$

as a more efficient implementation of  $\text{realizable}^{\text{GR1}}$ . Finally, we have

$$\text{cNot}_{\text{SD}}^{\text{GR1}}(B) = \begin{cases} \text{None} & \text{if } \text{realizable}_E^{\text{GR1}}((A, \exists(Y \cap B) . G \setminus B)) \\ \text{min}_{\text{SD}}^{\text{GR1}}((G \cup Y) \setminus B) & \text{otherwise,} \end{cases}$$

with  $\text{min}_{\text{SD}}^{\text{GR1}}(M) = \text{linMin}^{\text{test}_{\text{SD}}^2}(\text{ddmin}^{\text{test}_{\text{SD}}^1}(M))$ . The algorithm  $\text{ddmin}$  operates on  $\text{test}_{\text{SD}}^1$  defined as  $\text{test}_{\text{SD}}^1(M') = \mathbf{x} \Leftrightarrow \neg \text{realizable}_O^{\text{GR1}}((A, \exists(Y \setminus M') . G \cap M'))$ , and thus returns an over-approximation of the minimal unrealizable core. This over-approximation is further reduced by  $\text{linMin}$ , operating on  $\text{test}_{\text{SD}}^2(M') = \mathbf{x} \Leftrightarrow \text{realizable}_E^{\text{GR1}}((A, \exists(Y \setminus M') . G \cap M'))$ , and yielding the exact core. We use  $\text{linMin}$  instead of  $\text{ddmin}$  in the second step, because  $\text{ddmin}$  does not perform well when given an almost minimal set. It would waste many checks until the granularity is high enough to actually remove elements. As another performance optimization, we use early termination in the fixpoint computations. That is, if realizability or unrealizability is implied by an iterate of a fixpoint already, we abort the computation. Furthermore, for  $\text{realizable}_E^{\text{GR1}}$ , we use the set  $Q \setminus C_{\text{env}}^{\text{GR1}}$  of states as a starting point for the outermost greatest fixpoint computation of  $W_{\text{sys}}^{\text{GR1}}$ .

## 5 Experimental Results

We implemented the diagnosis approach for the class of GR(1) specifications inside Marduk, the back-end of RATSY [1]. The implementation<sup>2</sup> as well as the scripts<sup>3</sup> to reproduce our performance results are available for download. In this section we first give an example to illustrate the usefulness of our debugging approach. After that, we provide performance results for some more benchmarks.

### 5.1 Example

Our example specification (G5wst2 in Table 1) describes a generalized buffer connecting 5 senders to 2 receivers. It consists of 24 signals and 110 properties. We use lower case letters for inputs and upper case letters for outputs. The inputs `stob_req $i$` , with  $0 \leq i \leq 4$ , are used to signal that sender  $i$  requests to send data. With the outputs `BTOS_ACK $i$`  the buffer can acknowledge to sender  $i$  that sending is allowed. The buffer also communicates with a FIFO storage unit. The output `ENQ` is set if data should be enqueued into the FIFO. All other signals are not relevant for this example. A guarantee  $G(\text{ENQ}=0)$  has been added artificially to make the specification unrealizable. This guarantee forbids the buffer to enqueue data into the FIFO, which makes it impossible to handle any requests by the senders. An unrealizable core contains the system guarantees

$$\text{BTOS\_ACK4}=0 \wedge \text{ENQ}=0 \wedge \text{DEQ}=0 \tag{5}$$

$$G((\text{BTOS\_ACK4}=0 \wedge \text{X BTOS\_ACK4}=1) \Rightarrow \text{X ENQ}=1) \tag{6}$$

$$G(\text{ENQ}=0) \tag{7}$$

$$G F(\text{stob\_req4}=1 \Leftrightarrow \text{BTOS\_ACK4}=1) \tag{8}$$

$$G((\text{rtob\_ack0}=1 \wedge \text{X rtob\_ack0}=0) \Rightarrow \text{X DEQ}=1) \tag{9}$$

$$G((\text{rtob\_ack1}=1 \wedge \text{X rtob\_ack1}=0) \Rightarrow \text{X DEQ}=1) \tag{10}$$

$$G(\text{empty}=1 \Rightarrow \text{DEQ}=0) \tag{11}$$

and the outputs `BTOS_ACK4`, `ENQ`, and `DEQ`. This means that the system could not fulfill the listed guarantees even if it could set all outputs other than `BTOS_ACK4`, `ENQ`, and `DEQ` completely arbitrarily. This is the case because `BTOS_ACK4=0` initially (Eq. 5). It cannot change to 1 due to Eq. 6 and 7. If `stob_req4=1` forever, then the fairness guarantee stated in Eq. 8 cannot be fulfilled. The remaining guarantees (Eq. 9 to 11) are in the core because without them the system could enforce a violation of one environment assumption.

When the user is given this core (as suggested by [5]), or even an explanation [17] thereof, she has to face the following problem: She is presented only one inconsistency in the specification, but many more may exist, and she has to resolve all of them. There are often many ways to fix the presented conflict but only a few ways to fix all conflicts in the specification simultaneously. Our

<sup>2</sup> <http://rat.fbk.eu/ratsy/index.php/Main/Download>

<sup>3</sup> [http://www.iaik.tugraz.at/content/research/design\\_verification/others/](http://www.iaik.tugraz.at/content/research/design_verification/others/)

approach identifies guarantees and outputs which can be modified to resolve *all* conflicts. In our example, the user could learn from the presented core that Eq. 6 should be weakened. However, this does not make the original specification realizable, since other conflicts not containing Eq. 6 exist. The same argument can be made for Eq. 8. In fact, identical unrealizable cores exist with `BTOS_ACKi` instead of `BTOS_ACK4` for all  $0 \leq i \leq 3$ . What all conflicts have in common is the guarantee in Eq. 7 (the error we introduced). Our algorithm identifies exactly this guarantee and the signal `ENQ` as the only single-fault diagnoses. Hence, it provides much more accuracy in error localization than unrealizable cores or explanations thereof.

## 5.2 Performance Evaluation

We evaluated our diagnosis approach using the same benchmarks as in [16, 17]. These contain mutants of two parameterized specifications. One is an arbiter for the ARM AHB bus [2], parameterized with the number of masters it can handle. Mutants are denoted *Anei*, where *n* is the number of masters, *e* is the error that was introduced artificially, and *i* is a running index to distinguish different modifications of the same kind. The term `wœf` for *e* means that a fairness assumption was removed, `wsf` means that a fairness guarantee was added, and `wst` indicates that the specification was augmented with a safety guarantee. The second specification defines a generalized buffer [3], connecting *n* senders to two receivers. Mutants are denoted *Gnei* with the same syntax. All specification mutants have between 90 and 6004 properties and 22 to 218 signals. They are all satisfiable but unrealizable. The experiments were performed on an Intel Centrino 2 processor with  $2 \times 2.0$  GHz, 3 GB RAM, running 32-bit Linux.

Table 1 summarizes performance results. The Columns 1 to 3 give results for one minimal conflict (= unrealizable core) computation, using the algorithm of [17]. Column 1 gives the size of the conflict, Column 2 the number of guarantees in the conflict, and Column 3 the time needed for conflict computation. The Columns 4 to 8 summarize results for single-fault diagnoses computation. The number of single-fault diagnoses is given in Column 4. The next column lists the count of diagnoses that are guarantees. The time for diagnosis computation as described in Section 3.4 with `ddmin` as a minimization algorithm is shown in Columns 6. Column 7 presents the time for exactly the same computation, but when using the performance optimizations with approximations of realizability, as introduced in Section 4. Column 8 presents the according speed-up factor due to these optimizations. The Columns 9 to 12 finally give the number of diagnoses with at most two and three elements and the according computation times, respectively. Entries preceded by '>' indicate time-outs.

Our experimental results underline three statements. First, and most importantly, MBD gives a higher precision for error localization than the computation of a single unrealizable core. Single-fault diagnoses computation produces 40% less fault candidates than unrealizable core computation, where every unrealizable core element is taken as a fault candidate (Column 1 versus Column 4). That is, 40% of the unrealizable core elements cannot be modified in such a way

**Table 1.** Performance Results.

Column	One Conflict			Single-Fault Diagnoses					Mult. Fault Diagnoses						
	$ C $	$ C \cap G $	time (optimized)	$ \Delta_1 $	$ \Delta_1  = 1$	$ \Delta_1 $	$ \Delta_1 \subseteq G $	time	time, optimized	speed-up factor	$ \Delta_2 $	$ \Delta_2  \leq 2$	time: $\{\Delta_2\}$	$ \Delta_3 $	$ \Delta_3  \leq 3$
	[-]	[-]	[sec]	[-]	[-]	[sec]	[sec]	[-]			[-]	[sec]	[-]	[sec]	
A2woef1	9	4	0.5	7	2	2.6	1.5	1.8			14	13	47	48	
A4woef1	11	4	2.9	8	1	311	232	1.3			15	719	20	4.6 k	
A5woef1	12	4	12	9	1	2.5 k	960	2.6			16	6.1 k	21	>40 k	
A2wsf1	8	4	0.4	6	2	5.0	1.4	3.5			6	2.3	7	6.7	
A4wsf1	9	4	2.5	7	2	213	62	3.4			7	70	8	124	
A5wsf1	9	4	16	7	2	4.9 k	1.1 k	4.7			7	482	8	1.6 k	
A2wsf2	12	6	0.8	6	2	7.4	1.7	4.3			32	9.9	36	12	
A4wsf2	20	10	5.6	7	2	444	59	7.6			64	831	72	1.1 k	
A5wsf2	12	5	46	12	5	8.0 k	1.7 k	4.7			12	1.5 k	12	1.7 k	
A2wst1	9	4	0.5	8	3	3.6	1.3	2.8			8	1.7	12	3.8	
A4wst1	11	4	1.5	9	2	93	47	2.0			9	50	10	52	
A5wst1	12	4	5.0	10	2	1.6 k	336	4.8			10	348	10	365	
A2wst2	10	5	0.6	8	3	3.4	1.6	2.1			8	2.6	12	5.6	
A4wst2	12	5	2.0	9	2	115	50	2.3			9	55	10	62	
A5wst2	13	5	7.0	10	2	449	221	2.0			10	231	10	254	
G5woef1	15	8	1.4	10	6	3.3	3.0	1.1			14	7.2	27	20	
G20woef1	15	8	8.2	10	4	24	23	1.0			14	43	42	664	
G100woef1	15	8	125	10	4	1.2 k	1143	1.1			14	1.9 k	-	>40 k	
G5wsf1	19	11	3.5	14	7	6.2	6.1	1.0			24	12	24	19	
G20wsf1	49	26	929	44	22	787	1.0 k	0.8			54	1.2 k	54	1.2 k	
G100wsf1	-	-	>40 k	-	-	>40 k	>40 k	-			-	>40 k	-	>40 k	
G5wsf2	7	4	0.5	2	1	1.1	1.1	1.0			27	10	37	64	
G20wsf2	14	9	2.7	2	1	8.1	9.2	0.9			72	2.9 k	-	>40 k	
G100wsf2	7	4	116	2	1	404	407	1.0			-	>40 k	-	>40 k	
G5wst1	7	3	0.4	1	1	1.1	0.7	1.5			10	2.5	19	7.2	
G20wst1	7	3	2.1	1	1	6.3	5.3	1.2			10	18	10	42	
G100wst1	7	3	112	1	1	305	302	1.0			10	1.3 k	10	3.0 k	
G5wst2	9	6	0.8	2	1	1.6	1.3	1.2			8	4.2	8	14	
G20wst2	9	6	3.1	2	1	8.2	6.9	1.2			8	24	8	108	
G100wst2	9	6	120	2	1	402	368	1.1			8	1.2 k	8	4.2 k	
total	358	177	1.5 k	226	85	22 k	8.0 k	2.7			500	19 k	542	19 k	

that the specification becomes realizable. When only guarantees and no signals are considered, even more than half of the core elements can be excluded from being single-fault candidates (Column 2 versus Column 5). For `Gnwsf2`, `Gnwst1`, and `Gnwst2` only the one guarantee (and in two cases one signal of this guarantee) that was injected in order to make the specification unrealizable could be identified by our diagnosis approach. Also for `A4woef1` and `A5woef1`, only one guarantee is reported as modifiable in order to make the specification realizable. Moreover, in many cases, the number of diagnoses does not increase much with increasing bound on the cardinality of the diagnoses (Columns 9 and 11 versus Column 4). Hence, with MBD, the user often obtains very precise information stating where the unrealizability of the specification at hand can be fixed.

Second, MBD is more expensive than the computation of a single unrealizable core, unfortunately. Single-fault diagnosis computation takes about 5 times longer than unrealizable core computation (Column 3 versus Column 7). Without our performance optimizations, the situation is even worse (Column 3 versus Column 6). This is somewhat surprising, since there is only one additional realizability check per unrealizable core element required (for checking whether the core element is a single-fault diagnosis). The explanation for the large difference in the computation time is that most checks during unrealizable core computation are performed on small subsets of the original specification. In contrast, a check whether a given component forms a diagnosis is performed on nearly the entire specification, and is thus also much more expensive. For multiple-fault diagnosis computation, we can observe that the computation time increases noticeably with increasing bound on the cardinality of the diagnoses (Column 7 versus Columns 10 and 12). However, this is not considered as severe problem since the user is typically interested in small diagnoses anyway.

Third, the performance optimizations discussed in Section 4 seem to be effective. Using approximations of realizability, we obtain a speed-up of factor 2.7 when compared to the straightforward implementation described in Section 3.4 (Column 6 versus Column 7, contrasted in Column 8).

## 6 Conclusion

In this work we showed how model-based diagnosis can be used to locate bugs in unrealizable specifications, although the setting is quite different to the classical model-based diagnosis setting. Our approach computes signals and properties that can be weakened in order to make the specification realizable. It also yields conflicts which can be explained using counterstrategies to gain deeper insight into the problem. Hence, this work complements [17] in a nice way. Although our diagnosis approach is certainly a valuable debugging aid on its own, we believe that the combination with [17], as shown in Fig. 1, provides the user with an even more powerful debugging tool.

Experimental results for GR(1) specifications are promising, especially for single-fault diagnoses. Compared to unrealizable core computation, model-based diagnosis produces 40% less fault candidates. However, this improvement in the

accuracy comes at the price of a higher computational effort. Hence, model-based diagnosis can be considered as a stronger weapon against unrealizability. It may be an overkill for simpler bugs, where a glimpse on an unrealizable core suffices to fix the problem. However, the higher precision for error localization can be very important for more trickier bugs involving many properties. In order to tackle the performance problem, we showed on the example of GR(1) formulas how realizability and unrealizable cores can be computed faster using approximations. We achieve a speed-up factor of 2.7 for single-fault diagnoses computation. The concept is not specific to GR(1) and can be used for other logics as well.

In the future, we plan to investigate techniques to rule out diagnoses with additional user input, e.g., simulation traces. Furthermore, we plan to experiment with more fine-grained component definitions such as parts of properties.

## References

1. R. Bloem, A. Cimatti, K. Greimel, G. Hofferek, R. Könighofer, M. Roveri, V. Schuppan, and R. Seeber. RATSU - a new requirements analysis tool with synthesis. In *CAV*, pages 425–429. Springer, 2010. LNCS 6174.
2. R. Bloem, S. Galler, B. Jobstmann, N. Piterman, A. Pnueli, and M. Weiglhofer. Automatic hardware synthesis from specifications: A case study. In *DATE*, pages 1188–1193. ACM, 2007.
3. R. Bloem, S. Galler, B. Jobstmann, N. Piterman, A. Pnueli, and M. Weiglhofer. Specify, compile, run: Hardware from PSL. *Electronic Notes in Theoretical Computer Science*, 190(4):3–16, 2007.
4. H. Chockler, O. Kupferman, R. P. Kurshan, and M. Y. Vardi. A practical approach to coverage in model checking. In *CAV*, pages 66–78. Springer, 2001. LNCS 2102.
5. A. Cimatti, M. Roveri, V. Schuppan, and A. Tchaltev. Diagnostic information for realizability. In *VMCAI*, pages 52–67. Springer, 2008. LNCS 4905.
6. K. Claessen. A coverage analysis for safety property lists. In *FMCAD*, pages 139–145. IEEE, 2007.
7. L. Console, G. Friedrich, and D. Theseider Dupré. Model-based diagnosis meets error diagnosis in logic programs. In *IJCAI*, pages 1494–1499. Morgan-Kaufmann, 1993.
8. S. Das, A. Banerjee, P. Basu, P. Dasgupta, P. P. Chakrabarti, C. R. Mohan, and L. Fix. Formal methods for analyzing the completeness of an assertion suite against a high-level fault model. In *VLSI Design*, pages 201–206. IEEE, 2005.
9. J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.
10. S. Dellacherie. Automatic bus-protocol verification using assertions. In *Global Signal Processing Expo Conference (GSPx)*, 2004.
11. E. Filiot, N. Jin, and J.-F. Raskin. An antichain algorithm for LTL realizability. In *CAV*, pages 263–277. Springer, 2009. LNCS 5643.
12. D. Fisman, O. Kupferman, S. Seinfeld, and M. Y. Vardi. A framework for inverter vacuity. In *HVC*. Springer, 2008. LNCS 5394.
13. Y. V. Hoskote, T. Kam, P.-H. Ho, and X. Zhao. Coverage estimation for symbolic model checking. In *DAC*, pages 300–305, 1999.
14. B. Jobstmann and R. Bloem. Optimizations for LTL synthesis. In *FMCAD*, pages 117–124. IEEE, 2006.



15. S. Katz, O. Grumberg, and D. Geist. “Have I written enough properties?” — A method of comparison between specification and implementation. In *CHARME*, pages 280–297. Springer, 1999. LNCS 1703.
16. R. Könighofer. Debugging formal specifications with simplified counterstrategies. Master’s thesis, IAIK, Graz University of Technology, Inffeldgasse 16a, A-8010 Graz, Austria, 2009.
17. R. Könighofer, G. Hofferek, and R. Bloem. Debugging formal specifications using simple counterstrategies. In *FMCAD*, pages 152–159. IEEE, 2009.
18. D. Kozen. Results on the propositional  $\mu$ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
19. O. Kupferman and M. Y. Vardi. Safrless decision procedures. In *FOCS*, pages 531–542. IEEE, 2005.
20. A. Morgenstern and K. Schneider. Exploiting the temporal logic hierarchy and the non-confluence property for efficient LTL synthesis. *CoRR*, abs/1006.1408, 2010.
21. B. Peischl and F. Wotawa. Automated source-level error localization in hardware designs. *IEEE Design and Test of Computers*, 23:8–19, 2006.
22. I. Pill, S. Semprini, R. Cavada, M. Roveri, R. Bloem, and A. Cimatti. Formal analysis of hardware requirements. In *DAC*, pages 821–826. ACM, 2006.
23. N. Piterman, A. Pnueli, and Y. Sa’ar. Synthesis of reactive(1) designs. In *VMCAI*, pages 364–380. Springer, 2006. LNCS 3855.
24. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989.
25. R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, 1987.
26. M. Stumptner and F. Wotawa. Debugging functional programs. In *IJCAI*, pages 1074–1079. Morgan Kaufmann, 1999.
27. F. Wotawa. Debugging VHDL designs using model-based reasoning. *Artificial Intelligence in Engineering*, 14(4):331–351, 2000.
28. A. Zeller and R. Hildebrandt. Simplifying and isolating failure-inducing input. *IEEE Transactions on Software Engineering*, 28(2):183–200, 2002.