

# Machine-Checked Proofs For Realizability Checking Algorithms

Andreas Katis<sup>1</sup>, Andrew Gacek<sup>2</sup>, and Michael W. Whalen<sup>3</sup>

<sup>1</sup> <sup>3</sup>Department of Computer Science and Engineering  
University of Minnesota, 200 Union Street, Minneapolis, MN 55455, USA  
katis001@umn.edu,whalen@cs.umn.edu

<sup>2</sup> Rockwell Collins Advanced Technology Center  
400 Collins Road NE, Cedar Rapids, IA, 52498, USA  
andrew.gacek@gmail.com

**Abstract.** *Virtual integration* techniques focus on building architectural models of systems that can be analyzed early in the design cycle to try to lower cost, reduce risk, and improve quality of complex embedded systems. Given appropriate architectural descriptions, assume/guarantee contracts, and compositional reasoning rules, these techniques can be used to prove important safety properties about the architecture prior to system construction. For these proofs to be meaningful, each leaf-level component contract must be *realizable*; i.e., it is possible to construct a component such that for any input allowed by the contract assumptions, there is some output value that the component can produce that satisfies the contract guarantees.

We have recently proposed (in [1]) a contract-based realizability checking algorithm for assume/guarantee contracts over infinite theories supported by SMT solvers such as linear integer/real arithmetic and uninterpreted functions. In that work, we used an SMT solver and an algorithm similar to k-induction to establish the realizability of a contract, and justified our approach via a hand proof. Given the central importance of realizability to our virtual integration approach, we wanted additional confidence that our approach was sound. This paper describes a complete formalization of the approach in the Coq proof and specification language. During formalization, we found several small mistakes and missing assumptions in our reasoning. Although these did not compromise the correctness of the algorithm used in the checking tools, they point to the value of machine-checked formalization. In addition, we believe this is the first machine-checked formalization for a realizability algorithm.

## 1 Introduction

An ongoing effort at Rockwell Collins and The University of Minnesota has explored algorithms and tools for compositional proofs of correctness. The idea is to support hierarchical design and analysis of complex system architectures and co-evolution of requirements and architectures at multiple levels of abstraction [2]. We have created the AGREE reasoning framework [3] to support compositional

assume/guarantee contract reasoning over system architectural models written in AADL.

The soundness of the compositional argument requires that each leaf-level component contract is *realizable*; i.e., it is possible to construct a component such that for any input allowed by the contract assumptions, there is some output value that the component can produce that satisfies the contract guarantees. Unfortunately, without engineering support it is all too easy to write contracts of leaf-level components that can't be realized. When applying our tools in both industrial and classroom settings, this issue has led to incorrect compositional “proofs” of systems; in fact the goal of producing a compositional proof can lead to engineers modifying component-level requirements such that they are no longer possible to implement. In order to make our approach reasonable for practicing engineers, tool support must be provided for checking realizability.

The notion of realizability has been well-studied for many years [4,5,6,7,8,9], both for component synthesis and checking correctness of propositional temporal logic requirements. Checking realizability for contracts involving theories, on the other hand, is still an open problem. In recent work [1], we described a new approach for checking realizability of contracts as a Satisfiability Modulo Theories (SMT) problem and demonstrated its usefulness on several examples. Our approach is similar to k-induction [10] over quantified formulas. In that work, we provided hand-proofs for several aspects of two algorithms related to the soundness of the approach with respect to both proofs and counterexamples.

Unfortunately, hand proofs of complex systems often contain errors. Given the criticality of realizability checking to our tool chain and the soundness of our computational proofs, we would like a higher level of assurance than hand proofs can provide. In this paper, we provide a formalization of machine-checked proofs of correctness that ensure that the proposed realizability algorithms will perform as expected, using the Coq proof assistant.<sup>3</sup> The facilities in Coq, notably mixed use of induction and co-induction, make the construction of the proofs relatively straightforward. This approach illustrates how interactive theorem proving and SMT solving can be used together in a profitable way. Interactive theorem proving is used for describing the soundness of the checking algorithm (described in this paper). The algorithm is then implemented using a SMT solver, which can automatically solve complex verification instances.

The main contribution of this paper is, therefore, the first machine-checked formalization (to our knowledge) of a realizability checking algorithm. This is an important problem for both compositional verification involving virtual integration and component synthesis. In addition, the formalization process exposed errors regarding our initial definitions, including necessary assumptions to one of the main theorems to be proved and an error in the definition of realizability itself. While these errors did not ultimately impact the correctness of the algorithm, they underscore the importance of machine-checked proof.

---

<sup>3</sup> The Coq file is available at <https://github.com/andrewkatis/Coq/blob/master/realizability/Realizability.v>

In Section 2 we provide information on the Coq proof assistant. Section 3 contains the necessary informal background towards understanding our realizability checking approach. Sections 4.1 and 4.2 describe the definitions and theorems that were used both for defining realizability and the algorithms. In Section 5 we provide details on the algorithm’s implementation. Finally, in Section 6 we discuss our experience from the process of defining realizability and the various changes that were made along the way, and we report our conclusions in Section 7.

## 2 The Coq Proof Assistant

Coq<sup>4</sup> is an interactive tool used to formalize mathematical expressions and algorithms, and prove theorems regarding their correctness and functionality [11]. The tool was a result of the work on the calculus of constructions [12]. Its uses in the context of computer science vary, such as being a tool to represent the structure of a programming language and its characteristics, as well as to prove the correctness of underlying procedures in compilers. Compared to other mainstream interactive theorem provers, Coq is a tool that provides support on several aspects, such as the use of dependent types, as opposed to the Isabelle theorem prover [13], and proof by reflection, which is not supported by the PVS proof assistant [14]. A particularly essential feature is the tool’s support for inductive and coinductive definitions. Definitions using the *Inductive* type in Coq represent a least fixpoint of the corresponding type and are always accompanied by an induction principle, which is implicitly used to progress through a proof by applying induction on the definition. *CoInductive* definitions, on the other hand, represent a greatest fixpoint to their type. They describe a set containing every finite or infinite instance of that type, and their proofs are essentially infinite processes, built in a one-step fashion and requiring the existence of a guard condition that needs to hold for them to remain well-formed. Coinductive definitions allow a natural expression of infinite traces, which are central to our formalization of realizability, and are tedious to prove with hand-written proofs.

## 3 Realizability Checking

In [1] we presented our approach to the problem of realizability checking, introducing an algorithm involving the use of theories, a concept that, to the best of our knowledge, has yet to be examined. The realizability checks are defined over *assume-guarantee contracts*. Informally, *assumptions* describe the expectations of the component on its environment, usually in terms of component inputs. The *guarantees* describe the properties that will hold with respect to component outputs given that the assumptions are met. A contract *holds* on an infinite trace if either the assumption is violated or the guarantee holds throughout the trace.

---

<sup>4</sup> The Coq Proof Assistant is available at <https://coq.inria.fr/>

To illustrate, consider a system with a single integer input  $in$  and output  $out$  and a contract consisting of no assumptions and two guarantees:  $out = 2 * in$  and  $out \geq 0$ . This contract is not realizable. At issue is the behavior of the system if  $in < 0$ . In this case, the output of the system must both be positive and equal to  $2 * in$ , which is not possible. While this example is trivial, it can be very difficult to determine whether a contract involving dozens or hundreds of assumptions and guarantees is realizable. In [1], we describe two large-scale compositional reasoning examples (one medical device and one flight control system) that contained unrealizable leaf-level contracts that were previously unknown that were detected by our tools.

Informally, a *realizable contract* is one for which there exists a *transition system* that correctly and completely implements the contract. By “correctly” we mean that the transition system always produces outputs that satisfy the guarantees as long as the assumptions have always been met, and by “completely” we mean that the transition system never deadlocks on an input, so long as the assumptions have always been met. We will make these definitions precise in the next section.

This definition, while providing the proper theoretical basis for realizability, is not actually useful for constructing our checking algorithm. At issue is that our current algorithm provides no way to construct this ‘witness’ transition system (doing so would solve the general problem of program synthesis over contracts with theories, which we are currently researching). We therefore propose an alternative definition, according to which a contract is realizable if there exists a *viable path* consisting of *viable states*. A viable state is one where, for any inputs that satisfy the assumptions, there are outputs that satisfy the guarantees and lead to another viable state. This alternative definition requires that the contract be able to start in a viable state.

To derive checking algorithms from first principles, we first demonstrate that the two definitions (transition systems and viability) are equivalent. We can then use the viable definition as the basis of an algorithm for realizability checking. This algorithm consists of a *base check*, which ensures that there exists a finitely viable state for paths of length at least  $n$ , and an *extend check* to show that all the valid paths can be further extended in response to any input. Unfortunately, the complexity of the base check does not allow for an SMT solver to handle it efficiently. Because of this, we propose a simplified version of the algorithm including a base check that ensures the extendability of every valid path consisting of viable states. This check is only guaranteed sound with regard to ‘realizable’ results, that is, it may generate “false positives” in which the tool declares a contract unrealizable when in fact it can be realized. In early experiments, however, the tool results have been accurate.

## 4 Formalization in Coq

In the next two subsections, we will describe the formalization and proofs of these ideas in Coq. Section 4.1 will describe the definitions of realizability, while

Section 4.2 will describe the algorithms for realizability checking and their proofs of adequacy with respect to the definitions. To provide a graphical overview of the proof process, Figure 1 describes the connections between the various definitions, lemmas, and theorems in our work.

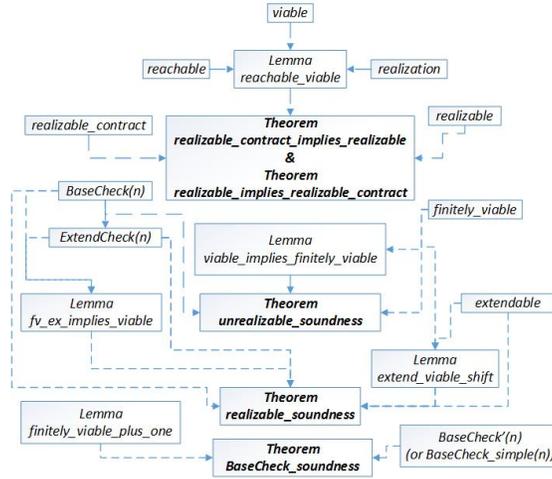


Fig. 1. Proof Graph

#### 4.1 Definitions

The types *state* and *inputs* are used to represent a state, and a given set of inputs. We use Coq’s *Prop* definition to describe the logical propositions regarding the component’s *transition system* through a set *I* of *initial* states and the *transition* relation *T* between two states and a set of inputs. Finally, the contract is defined by its *assumption* and *guarantee*, with the latter being implicitly referenced by a pair of initial and transitional guarantees (*iguarantee* and *tguarantee*). The corresponding definitions in Coq are shown below. Note that we do not expect that a contract would be defined over all variables in the transition system – rather its outputs – but we do not make any distinction between internal state variables and outputs in the formalism. This way, we can use state variables to, in some cases, simplify statements of guarantees.

- Inductive *inputs* : Type :=  
 $input : id \rightarrow nat \rightarrow inputs.$
- Inductive *state* : Type :=  
 $st : id \rightarrow nat \rightarrow state.$
- Definition *initial* :=  $state \rightarrow Prop.$

- Definition  $transition := state \rightarrow inputs \rightarrow state \rightarrow Prop$ .
- Definition  $iguarantee := state \rightarrow Prop$ .
- Definition  $tguarantee := state \rightarrow inputs \rightarrow state \rightarrow Prop$ .
- Definition  $assumption := state \rightarrow inputs \rightarrow Prop$ .

A state  $s$  is *reachable* with respect to the given assumptions if there exists a path from an initial state to  $s$ , while each transition in the path is satisfying the assumptions. Given a contract  $(A, (G_I, G_T))$ , a transition system  $(I, T)$  is its *realization* if the following four conditions hold:

1.  $\forall s. I(s) \Rightarrow G_I(s)$
2.  $\forall s, i, s'. reachable_A(s) \wedge A(s, i) \wedge T(s, i, s') \Rightarrow G_T(s, i, s')$
3.  $\exists s. I(s)$
4.  $\forall s, i. reachable_A(s) \wedge A(s, i) \Rightarrow \exists s'. T(s, i, s')$

Finally, we define that a given contract is *realizable*, if the existence of a transition system, which is a realization of the contract, is proved. The formalized definitions in Coq for the *reachable* state, the *realization* of a contract and whether it is *realizable* follow.

- Inductive  $reachable (s : state) (I : initial) (T : transition) (A : assumption) : Prop :=$   
 $rch :$   
 $((I s) \vee$   
 $((\exists (s' : state) (inp : inputs),$   
 $(reachable s' I T A) \wedge (A s' inp) \wedge (T s' inp s))) \rightarrow$   
 $reachable s I T A).$
- Inductive  $realization (I : initial) (T : transition) (A : assumption) (G_I : iguarantee) (G_T : tguarantee) : Prop :=$   
 $real : ((\forall (s : state), (I s) \rightarrow (G_I s)) \wedge$   
 $(\forall (s s' : state) (inp : inputs),$   
 $((reachable s I T A) \wedge (A s inp) \wedge (T s inp s')) \rightarrow G_T s inp s') \wedge$   
 $(\exists (s : state), I s) \wedge$   
 $(\forall (s : state) (inp : inputs), (reachable s I T A \wedge (A s inp)) \rightarrow$   
 $(\exists (s' : state), T s inp s')))) \rightarrow$   
 $realization I T A G_I G_T.$
- Inductive  $realizable_contract (A : assumption) (G_I : iguarantee) (G_T : tguarantee) : Prop :=$   
 $rc : (\exists (I : initial) (T : transition), realization I T A G_I G_T) \rightarrow$   
 $realizable_contract A G_I G_T.$

While the definitions of *realization* and *realizable\_contract* are quite straightforward, they cannot be used directly to construct an actual realizability checking algorithm. Therefore, we proposed the notion of a state being *viable* with

respect to a contract, meaning that the transition system continues to be a realization of the contract, while we are at such a state. In other words, a state is *viable* ( $viable(s)$ ) if the transitional guarantee  $G_T$  infinitely holds, given valid inputs. Using the definition of *viable*, a contract is *realizable* if and only if  $\exists s. G_I(s) \wedge viable(s)$ .

- CoInductive  $viable (s : state) (A : assumption) (G_I : iguarantee) (G_T : tguarantee) : Prop :=$   
 $vbl : (\forall (inp : inputs), (A s inp) \rightarrow$   
 $(\exists (s' : state), G_T s inp s' \wedge viable s' A G_I G_T)) \rightarrow$   
 $viable s A G_I G_T.$
- Inductive  $realizable (A : assumption) (G_I : iguarantee) (G_T : tguarantee) : Prop :=$   
 $rl : (\exists (s : state), G_I s \wedge viable s A G_I G_T) \rightarrow realizable A G_I G_T.$

Having a more useful definition for realizability, we need to prove the equivalence between the definitions of *realizable\_contract* and *realizable*. The Coq definition of the theorem was split into two separate theorems, each for one of the two directions of the proof. Towards the two proofs, the auxiliary lemma that, given a realization,  $\forall s. reachable_A(s) \Rightarrow viable(s)$  is necessary.

- Lemma  $reachable\_viable : \forall (s : state) (I : initial) (T : transition) (A : assumption) (G_I : iguarantee) (G_T : tguarantee),$   
 $realization I T A G_I G_T \rightarrow reachable s I T A \rightarrow viable s A G_I G_T.$

The informal proof of the lemma relies initially on the unrolling of the *viable* definition, for a specific state  $s$ . Thus, we are left to prove that there exists another state  $s'$  that we can traverse into, in addition to being viable. The former can be proved directly from the conditions 2 and 4 of the definition of *realization*. For the latter, by the definition of *viable* on  $s'$  we need to show that  $s'$  is reachable. Given the definition of *reachable* though, we just need to prove that there exists another reachable state from which we can reach  $s'$ , in one step. But we already know that  $s$  is such a state, and thus the lemma holds.

- Theorem  $realizable\_contract\_implies\_realizable (I : initial) (T : transition) : \forall (A : assumption) (G_I : iguarantee) (G_T : tguarantee),$   
 $realizable\_contract A G_I G_T \rightarrow realizable A G_I G_T.$
- Theorem  $realizable\_implies\_realizable\_contract (I : initial) (T : transition) : \forall (A : assumption) (G_I : iguarantee) (G_T : tguarantee),$   
 $realizable A G_I G_T \rightarrow realizable\_contract A G_I G_T.$

The first part of the theorem requires us to prove that there exists a viable state  $s$  for which the initial guarantee holds. Considering that we have a contract that is realizable under the *realizable\_contract* definition, we have a transition system that is a realization of the contract, and thus from the third condition of the *realization* definition, there exists an initial state  $s'$  for which, using the

first condition, the initial guarantee holds. Thus, we are left to prove that  $s'$  is viable. But, by proving that  $s'$  is reachable, we can use the *reachable\_viable* lemma to show that  $s'$  is indeed viable.

The second direction requires a bit more effort. Assuming that we have a viable state  $s_0$  with  $G_I(s_0)$  being true, we define  $I(s) = (s = s_0)$  and  $T(s, inp, s') = G_T(s, inp, s') \wedge viable(s')$ . Initially, we need to prove the *reachable\_viable* lemma in this context, with the additional assumption that another viable state already exists ( $s_0$  in this case). Having done so, we need to prove that there exists a transition system that is a realization of the given contract. Given the transition system that we defined earlier, we need to show that each of the four conditions hold. Since  $I(s) = (s = s_0)$  and  $G_I(s_0)$  hold, the proof for the first condition is trivial. Using the assumption that  $T(s, inp, s') = G_T(s, inp, s') \wedge viable(s')$ , we can also trivially prove the second condition, while the third condition is simply proved by reflexivity on the state  $s_0$ . Finally, for the fourth condition we need to prove that  $\forall s, inp. reachable_A(s) \wedge A(s, inp) \Rightarrow \exists s'. G_T(s, inp, s') \wedge viable(s')$ . By applying the *reachable\_viable* lemma on the reachable state  $s$  in the assumptions, we show that  $s$  is also viable, if  $s_0$  is viable, which is what we assumed in the first place. Thus, coming back into what we need to prove, and unrolling the definition of *viable* on  $s$ , we have that  $\forall inp. A(s, inp) \Rightarrow \exists s'. G_T(s, inp, s') \wedge viable(s')$  which completes the proof.

## 4.2 Algorithms

In this section we provide a description of the formalization and proof of soundness of our realizability checking algorithms. Initially, we define an under-approximation of the definition of viability, for the finite case. Thus, a state is *finitely\_viable* for  $n$  steps (*viable<sub>n</sub>*( $s$ )), if the transitional guarantee  $G_T$  holds for at least  $n$  steps, given valid inputs.

- Inductive *finitely\_viable* :  $nat \rightarrow state \rightarrow assumption \rightarrow tguarantee \rightarrow Prop$  :=
  - | *fnil* :  $\forall s A G_T, finitely\_viable\ 0\ s\ A\ G_T$
  - | *fv* :  $\forall n\ s\ A\ G_T, finitely\_viable\ n\ s\ A\ G_T \rightarrow$   
 $(\forall (inp : inputs), A\ s\ inp \rightarrow (\exists s', G_T\ s\ inp\ s')) \rightarrow$   
 $finitely\_viable\ (S\ n)\ s\ A\ G_T.$

In addition to the *finitely\_viable* definition, an under-approximation of viability is also used, called one-step extension. Therefore, a valid path leading to a state  $s$  is *extendable* after  $n$  steps, if any path from  $s$ , of length at least  $n$ , can be further extended given a valid input.

- Inductive *extendable* :  $nat \rightarrow state \rightarrow assumption \rightarrow tguarantee \rightarrow Prop$  :=
  - | *exnil* :  $\forall (s : state) (A : assumption) (G_T : tguarantee),$   
 $(\forall (inp : inputs), A\ s\ inp \rightarrow \exists (s' : state), G_T\ s\ inp\ s') \rightarrow$   
 $extendable\ 0\ s\ A\ G_T$

$$\begin{aligned}
& | \text{ex} : \forall n \ s \ A \ G_T, \\
& \quad (\forall \text{inp } s', A \ s \ \text{inp} \wedge G_T \ s \ \text{inp } s' \wedge \text{extendable } n \ s' \ A \ G_T) \rightarrow \\
& \quad \text{extendable } (S \ n) \ s \ A \ G_T.
\end{aligned}$$

**An Exact Algorithm for Realizability Checking** The algorithm that we propose for realizability checking consists of two checks. The *BaseCheck*( $n$ ) procedure ensures that  $\exists s. G_I(s) \wedge \text{viable}_n(s)$ , while *ExtendCheck*( $n$ ) makes sure that the given state from *BaseCheck* is extendable for any  $n$ .

- Definition *BaseCheck* ( $n : \text{nat}$ ) ( $A : \text{assumption}$ ) ( $G_I : \text{iguarantee}$ ) ( $G_T : \text{tguarantee}$ ) :=  
 $\exists (s : \text{state}), (G_I \ s \wedge \text{finitely\_viable } n \ s \ A \ G_T)$ .
- Definition *ExtendCheck* ( $n : \text{nat}$ ) ( $A : \text{assumption}$ ) ( $G_T : \text{tguarantee}$ ) :=  
 $\forall s \ A \ G_T, \text{extendable } n \ s \ A \ G_T$ .

Using the *BaseCheck*( $n$ ) and *ExtendCheck*( $n$ ), the algorithm determines the realizability of the given contract, using the following procedure.

```

for  $n = 0$  to  $\infty$  do
  if not BaseCheck( $n$ ) then
    return "unrealizable"
  else if ExtendCheck( $n$ ) then
    return "realizable"
  end if
end for

```

Using the definitions of *BaseCheck* and *ExtendCheck*, we proved the algorithm's soundness, both for the 'unrealizable' and 'realizable' case. The main idea behind the proof of soundness for the 'unrealizable' result is to prove the contrapositive, that is, given a realizable contract, there exists a natural number  $x$  for which *BaseCheck*( $x$ ) holds. Unfolding the definition of *BaseCheck*( $x$ ), we need to show that  $\exists s. G_I(s) \wedge \text{viable}_x(s)$ . Knowing that our assumption *realizable\_contract*  $A \ G_I \ G_T$  is equivalent to the *realizable* definition, provides us with a state  $s'$ , for which  $G_I(s') \wedge \text{viable}(s')$  holds. Here, we need an additional lemma, according to which  $\forall s, n. \text{viable}(s) \Rightarrow \text{viable}_n(s)$  (stated as *viable\_implies\_finitely\_viable* below). Thus, using the lemma on *viable*( $s'$ ) with  $n = x$ , we get that *viable* <sub>$x$</sub> ( $s'$ ), thus completing the proof.

- Lemma *viable\_implies\_finitely\_viable* :  $\forall s \ A \ G_I \ G_T \ n,$   
 $\text{viable } s \ A \ G_I \ G_T \rightarrow \text{finitely\_viable } n \ s \ A \ G_T$ .
- Theorem *unrealizable\_soundness* :  $\forall (I : \text{initial}) (T : \text{transition}) (A : \text{assumption}) (G_I : \text{iguarantee}) (G_T : \text{tguarantee}),$   
 $(\exists n, \neg \text{BaseCheck } n \ A \ G_I \ G_T) \rightarrow \neg \text{realizable\_contract } A \ G_I \ G_T$ .

For the soundness of the 'realizable' result, we first need to prove two lemmas. Initially, *extend\_viable\_shift*, shows the way that *Extend* <sub>$n$</sub> ( $s$ ) can be used to

shift  $viablen(s)$  forward. The proof for this lemma is done by using induction on  $n$ . The base case is proved trivially, by unfolding the definitions of *extendable* and *finitely\_viable* in the assumptions. For the inductive case, we assume that the same state  $s$  is extendable and finitely viable for paths of length  $n + 1$ , and try to prove that there exists a finitely viable state  $s'$  for paths of length  $n + 1$ , to which we can traverse from  $s$ , with the contract guarantees still holding after the transition. By considering that  $s$  is extendable for paths of length  $n + 1$ , we can use it as that potentially existing state in the proof, requiring that we can transition from  $s$  to itself, with the transitional guarantees staying true, and  $s$  being finitely viable for paths of length  $n + 1$ . The former is true through the definition of *extendable*, while the second is an already given assumption by the inductive step.

- Lemma *extend\_viable\_shift* :  $\forall (s : state) (n : nat) (inp : inputs) (A : assumption) (G_I : iguarantee) (G_T : tguarantee),$   
 $(extendable\ n\ s\ A\ G_T \wedge finitely\_viable\ n\ s\ A\ G_T \wedge A\ s\ inp) \rightarrow$   
 $(\exists s', G_T\ s\ inp\ s' \wedge finitely\_viable\ n\ s'\ A\ G_T).$
- Lemma *fv\_ex\_implies\_viable* :  $\forall (s : state) (n : nat) (A : assumption) G_I$   
 $G_T,$   
 $(finitely\_viable\ n\ s\ A\ G_T \wedge ExtendCheck\ n\ A\ G_T) \rightarrow viable\ s\ A\ G_I\ G_T.$
- Theorem *realizable\_soundness* :  $\forall (I : initial) (T : transition) A\ G_I\ G_T,$   
 $(\exists n, (BaseCheck\ n\ A\ G_I\ G_T \wedge ExtendCheck\ n\ A\ G_T)) \rightarrow$   
 $realizable\_contract\ A\ G_I\ G_T.$

To prove the theorem, we try to prove the equivalent for the *realizable* definition instead. The existence of a state for which the initial guarantees hold is derived from the assumption that *BaseCheck* holds for a finitely viable state, while the proof that the same state is also viable comes from the use of the *fv\_ex\_implies\_viable* lemma, which is proved through the use of *extend\_viable\_shift*.

**An Approximate Algorithm for Realizability Checking** Following the definition of our approach, we noticed the problematic nature of *BaseCheck*( $n$ ) having  $2n$  quantifier alternations, which cannot be handled efficiently by an SMT solver. To that end, we proposed a simplified version of the *BaseCheck*( $n$ ) procedure, called *BaseCheck'*( $n$ ), stated as *BaseCheck\_simple* below.

- Definition *BaseCheck\_simple* ( $n : nat$ ) ( $A : assumption$ ) ( $G_I : iguarantee$ ) ( $G_T : tguarantee$ ) :=  $\forall s, (G_I\ s) \rightarrow extendable\ n\ s\ A\ G_T.$
- Lemma *finitely\_viable\_plus\_one* :  $\forall s\ n\ A\ (gi : iguarantee) (G_T : tguarantee)$   
 $(inp : inputs),$   
 $(extendable\ n\ s\ A\ G_T \wedge finitely\_viable\ n\ s\ A\ G_T) \rightarrow$   
 $finitely\_viable\ (S\ n)\ s\ A\ G_T.$

- Theorem *BaseCheck\_soundness* :  $\forall n A (G_I : iguarantee) (G_T : tguarantee)$   
 $(i : inputs),$   
 $((\exists s, G_I s) \wedge (\forall k, (k \leq n) \rightarrow BaseCheck\_simple\ k\ A\ G_I\ G_T)) \rightarrow$   
 $BaseCheck\ n\ A\ G_I\ G_T.$

The simplified *BaseCheck'(n)*, while being an easier instance for an SMT solver, is not sound for the 'unrealizable' case, falsely reporting some realizable contracts to not be so. Nevertheless, we proved the modified algorithm's soundness for the 'realizable' result, with the use of an auxiliary lemma.

The lemma, *finitely\_viable\_plus\_one* simply refers to the fact that an extendable and finitely viable state *s*, for a given number of steps *n*, is also finitely viable for *n + 1* steps. The proof is done by induction on *n*. The base case is trivially proved, by the definition of *finitely\_viable*, and the assumption that *s* is extendable. For the inductive case, we use the inductive hypothesis, which leaves us to prove the assumptions on a specific state *s*. The extendability is trivially shown since we already know that *s* is extendable for paths of length *n + 1*, with the same idea being applied to prove that *s* is finitely viable for *n*.

Finally, the proof of soundness for the 'realizable' result of the *BaseCheck'(n)* procedure is done by using induction on *n*. The base case is trivially true, using the fact that all paths of zero length are finitely viable. The inductive step then requires us to prove that *BaseCheck(n + 1)* holds. In order to do so, we need to construct the inductive hypothesis' assumption, as a separate assumption to the theorem's scope. By applying the inductive hypothesis to the newly created assumption, we have that *BaseCheck(n)* holds. By unrolling the definition of *BaseCheck(n)* and applying the lemma *finitely\_viable\_plus\_one* on the extracted state, say *x*, we finally prove that *x* is extendable through the definition of *BaseCheck'(n)*, completing the proof at the same time.

Figure 1 provides a simplified proof graph of all the necessary definitions and partially, for graph simplicity purposes, the way that they are used towards proving the lemmas and theorems stated in this paper.

## 5 Implementation

The algorithm is now an optional feature, namely JRealizability in JKind [15], a Java implementation of the KIND 2 model checker,<sup>5</sup> and supports models expressed using the Lustre language [16], which are a result of AGREE's translation process of contracts written in AADL. A typical process for checking models in the above environment starts from providing the corresponding Lustre program to JKind, which JRealizability uses to find a number *n*,  $n \geq 0$ , such that both *BaseCheck'(n)* and *ExtendCheck(n)* hold. Specifically, the model's variables and contract are being translated in the SMT-LIB2 format, followed by the construction of each check's corresponding query for the current value of *n*, in its negated form. The resulting SMT-LIB2 file is provided as input to the

<sup>5</sup> You can download the KIND model checker at <http://kind2-mc.github.io/kind2/>

Z3 SMT solver [17], which attempts to answer the given query. In the case that the negated formula is unsatisfiable, JRealizability returns a 'realizable' result. On the other hand, a satisfiable query implies that the model is unrealizable. Consequently, the tool requests a model, i.e. an instance of the contract's variables that reflects Z3's result, and proceeds to construct a counterexample that describes the exact cause of the contract's unrealizability. Finally, in those cases where the quantified query is too difficult for Z3 to solve, an 'unknown' result is reported, both by Z3 and JRealizability.

The implementation was used in [1] to verify the correctness of contracts in terms of realizability in three different case studies. The performance was very good for the concrete results, with the tool exceeding its predefined timeout value for the 'unknown' ones. False positive results (*BaseCheck'(n)*) were not found during this process, as every unrealizable contract was manually proved to be a result of conflicts in the provided assumptions or guarantees. A final remark is the fact that the most critical case studies already had an implementation that was supposed to work correctly. As such, the discovery of unrealizable contracts in these systems eventually required a total revision of the formalized requirements defined for each system, thus hindering the development process.

## 6 Discussion

While our work on realizability is based on simple definitions, formalizing them and refining the algorithms in Coq was non-trivial. Proving the lemmas and theorems using Coq helped us discover minor errors in our informal statements. For example, our proof of the one-way soundness theorem for the simplified *BaseCheck* in [1] lacks the necessary assumption that there exists a state for which the initial guarantees hold. Another example is that we forgot to include initial states in our definition of reachable states in the informal proof. The use of a mechanized theorem prover exposed some missing knowledge in the informal text, and helped us provide a more precise version of the theorem. Although these errors in the hand proofs did not lead to problems with our implementation, Coq improved both our theorems and proofs, and provided a very high level of assurance that our algorithm is correct.

## 7 Conclusion

The work in this paper was particularly important towards verifying our approach and learning more about the actual functionality of the algorithm. Interactive theorem provers like Coq provide the necessary support to define the notions and assertions while being able to effectively prove theorems in a far more convenient and reassuring way, in contrast to hand-written, informal proofs, especially when it comes down to tracking formulas containing alternating quantifiers. Furthermore, the procedure of proving the theorems in an interactive way with a tool allowed us to refine our definitions. Additionally, the time that was required was minimal when compared to the process of considering the informal

proofs and writing down our requirements in English. The most important outcome was the proof of correctness of our approach that enabled us to provide a complementary set of definitions and proofs, easily processed by an experienced Coq user.

To conclude, there is substantial additional work that could be performed in terms of fleshing out the formalisms used in the proofs for our particular implementation. For example, we could define the structure and types of inputs and outputs, and describe how transition systems are realized in the AGREE tool suite. However, the work that has been performed shows the soundness of the proof system and our algorithms with respect to proofs of realizability, allowing us to proceed with very high confidence as to the correctness of our approach.

**Acknowledgments.** This work was funded by DARPA and AFRL under contract 4504789784 (Secure Mathematically-Assured Composition of Control Models), and by NASA under contract NNA13AA21C (Compositional Verification of Flight Critical Systems), and by NSF under grant CNS-1035715 (Assuring the safety, security, and reliability of medical device cyber physical systems).

## References

1. A. Gacek, A. Katis, M. W. Whalen, J. Backes, and D. Cofer, "Towards Realizability Checking of Contracts using Theories," in *Proceedings of the 7th NASA Formal Methods Symposium (NFM 2015)*.
2. M. W. Whalen, A. Gacek, D. Cofer, A. Murugesan, M. P. Heimdahl, and S. Rayadurgam, "Your what is my how: Iteration and hierarchy in system design," *Software, IEEE*, vol. 30, no. 2, pp. 54–60, 2013.
3. D. D. Cofer, A. Gacek, S. P. Miller, M. W. Whalen, B. LaValley, and L. Sha, "Compositional verification of architectural models," in *Proceedings of the 4th NASA Formal Methods Symposium (NFM 2012)*, A. E. Goodloe and S. Person, Eds., vol. 7226. Berlin, Heidelberg: Springer-Verlag, April 2012, pp. 126–140.
4. A. Pnueli and R. Rosner, "On the Synthesis of a Reactive Module," *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages (POPL'89)*, pp. 179–190, 1989.
5. A. Bohy, V. Bruy re, E. Filiot, N. Jin, and J.-F. Raskin, "Acacia+, a tool for LTL Synthesis," *Proceedings of the 24th International Conference on Computer Aided Verification (CAV'12)*, pp. 652–657, 2012.
6. J. Hamza, B. Jobstmann, and V. Kuncak, "Synthesis for Regular Specifications over Unbounded Domains," *Proceedings of the 2010 Conference on Formal Methods in Computer-Aided Design*, pp. 101–109, 2010.
7. K. Chatterjee and T. A. Henzinger, "Assume-Guarantee Synthesis," *Proceedings of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'07)*, pp. 261–275, 2007.
8. C. A. Gunter, E. L. Gunter, M. Jackson, and P. Zave, "A Reference model for Requirements and Specifications," *IEEE Software*, vol. 17, no. 3, pp. 37–43, May 2000.

9. L. M. Patcas, M. Lawford, and T. Maibaum, "From system requirements to software requirements in the four-variable model," in *AUTOMATED VERIFICATION OF CRITICAL SYSTEMS (AVOCS) 2013*. Citeseer, 2014.
10. M. Sheeran, S. Singh, and G. Stålmarck, "Checking safety properties using induction and a sat-solver," in *Formal Methods in Computer-Aided Design*. Springer, 2000, pp. 127–144.
11. The Coq Development Team, *The Coq Proof Assistant Reference Manual*, 8th ed., INRIA, 2012-2014.
12. T. Coquand and G. Huet, "Constructions: A higher order proof system for mechanizing mathematics," in *EUROCAL'85*. Springer, 1985, pp. 151–184.
13. L. C. Paulson, "The foundation of a generic theorem prover," *Journal of Automated Reasoning*, vol. 5, no. 3, pp. 363–397, 1989.
14. S. Owre, J. M. Rushby, , and N. Shankar, "PVS: A prototype verification system," in *11th International Conference on Automated Deduction (CADE)*, ser. Lecture Notes in Artificial Intelligence, D. Kapur, Ed., vol. 607. Saratoga, NY: Springer-Verlag, jun 1992, pp. 748–752. [Online]. Available: <http://www.csl.sri.com/papers/cade92-pvs/>
15. A. Gacek, "JKind - a Java implementation of the KIND model checker," <https://github.com/agacek/jkind>, 2014.
16. N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud, "The synchronous data flow programming language lustre," *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1305–1320, 1991.
17. L. De Moura and N. Bjørner, "Z3: An efficient SMT solver," in *Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2008, pp. 337–340.