

Realizability Checking of Contracts with Kind 2

(Draft)

Daniel Larráz¹ and Cesare Tinelli¹

Department of Computer Science, The University of Iowa. USA

Abstract. We present a new feature of the open-source model checker Kind 2 which checks whether a 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. When the contract is proven unrealizable, it provides a deadlocking computation and a set of conflicting guarantees. This new feature can be used to detect flaws in component specifications and to ensure the correctness of Kind 2’s compositional proof arguments.

1 Introduction

Contract-based software development has long been a leading methodology for the construction of component-based reactive systems, embedded systems in particular. Contracts provide a mechanism for capturing the information needed to specify and reason about component-level properties at a desired level of abstraction. In this paradigm, each component is associated with a contract specifying its input-output behavior in terms of guarantees provided by the component when its environment satisfies certain given assumptions. Contracts are an effective way to establish boundaries between components and can be used efficiently to prove global properties about a system prior to its construction. Such proofs are built upon the premise that each leaf-level component contract in the system hierarchy 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. However, without engineering support it is all too easy to write leaf-level components that cannot be realized.

This report describes a new feature of the open-source model checker KIND 2 [3] which allows users to verify the realizability of contracts. KIND 2 is an SMT-based model checker for safety properties of finite- and infinite-state synchronous reactive systems. It takes as input models written in an extension of the Lustre language [6] that allows the specification of assume-guarantee-style contracts for system components. KIND 2’s contract language [2] is expressive enough to allow one to represent any (LTL) regular safety property by recasting it in terms of invariant properties. One of KIND 2’s distinguishing features is its support for modular and compositional analysis of hierarchical and multi-component systems. KIND 2 traverses the subsystem hierarchy bottom-up, analyzing each system component, and performing fine-grained abstraction and refinement of the sub-components.

The behavior of each component can be specified by providing either a set of equations that define the component’s output in terms of its input and internal state (a *low-level* specification), or an assume-guarantee contract (a *high-level* specification), or both. The syntax restrictions and semantics of the Lustre language ensure that every low-level specification of a component is *executable* in the sense that for each possible input for the component and each internal state there is a unique output and next state for the component to move to. When both specifications are provided,

```

1 type digit_range = subrange [0,9] of int;
2 const MAX_TIME = 60 * 9 + 59;
3
4 node imported Display_Control(
5   cancel: bool; incr: bool; decr: bool; baking: bool
6 )
7 returns (
8   left_digit: digit_range; middle_digit: digit_range; right_digit: digit_range;
9   minutes_to_cook: int
10 );

```

Fig. 1: Display_Control component

the low-level specification is expected to be a refinement of the high-level one. KIND 2 checks this by verifying that every execution that satisfies the former also satisfies the latter. Informally, we say that the set of equations *satisfy* the contract. However, in compositional reasoning, when only a contract is provided for a subcomponent, KIND 2 assumes the existence of a component satisfying the contract when checking the satisfaction of the top-level component requirements, which may lead to bogus compositional proof arguments when the subcomponent’s contract is unrealizable.

Example 1. We will use a simple model to illustrate the concepts and the functionality of KIND 2 introduced in this report. Suppose we want to design a component that controls the display of an oven. The oven has a panel with three buttons: a *cancel* button, an *increase* button, and a *decrease* button. The component that controls the display reads the three button inputs from the panel and the current mode of the oven (baking or not baking), and it sets three digit displays showing the current number of minutes to cook accordingly. The left-most digit corresponds to hours, the middle digit is tens of minutes, and the right digit is minutes.

Our model for the component is described in KIND 2’s input language in Figure 1, which defines (starting at line 4) the component’s interface, and Figure 2, which contains its contract. The interface includes three inputs, `cancel`, `incr`, and `decr`, one for each button input, and an additional input, `baking`, to indicate whether the oven is in baking mode or not (line 5). The component has one output for each of the three digit displays (line 8) whose values range between 0 and 9 (line 1). In addition, the component has an additional output that reports the total time of the displayed digits (line 9). Following a model-based design, we model an abstraction of `Display_Control` component instead of specifying a complete set of equations that fully determine the behavior of the whole component. KIND 2 allows the user to specify contracts for individual nodes, either as special Lustre comments added directly inside the node declaration, or as the instantiation of an external stand-alone contract that can be imported in the body of other contracts. The contract of `Display_Control`, included directly in the node (lines 1-29 of Figure 2), specifies the relationship between the value of `minutes_to_cook` and the three displayed digits (guarantees G1-G3) as well as the value of `minutes_to_cook` in reaction to different situations (guarantees G4-G9). For instance, guarantee G8 specifies that `minutes_to_cook` shall increase by one, when the oven is not cooking, if the `incr` button is pressed and the previous total time to cook is less than `MAX_TIME`. Moreover, when the oven is not baking, `minutes_to_cook` shall be zero if the `incr` button is pressed but the previous total time to cook is not less than `MAX_TIME`.

This specification is detailed enough to prove some interesting properties about the component. In order to do that, we can wrap up an instance of the `Display_Control` component in an *observer*

```

1 (*@contract
2   -- The following three definitions are based on the fact that
3   -- minutes_to_cook shall match the total time of the displayed digits
4   guarantee "G1: The left-most digit corresponds to hours"
5     left_digit = (minutes_to_cook div 60);
6   guarantee "G2: The middle digit corresponds to tens of minutes"
7     middle_digit = (minutes_to_cook mod 60) div 10;
8   guarantee "G3: The right digit corresponds to minutes"
9     right_digit = (minutes_to_cook mod 10);
10
11   var any_button_pressed: bool = incr or decr or cancel; -- auxiliary var
12
13   guarantee "G4: minutes_to_cook shall be initially zero"
14     minutes_to_cook = 0 -> true;
15   guarantee "G5: If the cancel button is pressed, minutes_to_cook shall be zero"
16     cancel => minutes_to_cook = 0;
17   guarantee "G6: When baking, minutes_to_cook shall remain the same or decrease"
18     true -> baking => minutes_to_cook <= pre minutes_to_cook;
19   guarantee "G7: When not baking, if not button is pressed, minutes_to_cook shall not change"
20     true -> (not baking and not any_button_pressed) => minutes_to_cook = pre minutes_to_cook;
21   guarantee "G8: When not baking, if incr is pressed, minutes_to_cook shall
22     increase by one if it was less than MAX_TIME or be zero otherwise"
23     true -> (not baking and incr) =>
24       (minutes_to_cook = if pre minutes_to_cook < MAX_TIME then pre minutes_to_cook + 1 else 0);
25   guarantee "G9: When not baking, if decr is pressed but not incr, minutes_to_cook shall
26     decrease by one if it was greater than 0 or be MAX_TIME otherwise"
27     true -> (not baking and not incr and decr) =>
28       (minutes_to_cook = if pre minutes_to_cook > 0 then pre minutes_to_cook - 1 else MAX_TIME);
29 *)

```

Fig. 2: KIND 2's contract for the Display_Control component

component, specify the properties we want to check as guarantees in the contract of the observer component, and ask KIND 2 to check the satisfaction of the contract. For example, the contract of `Display_Control_Observer` in Figure 3 specifies three properties (P1-P3) about the behavior of `Display_Control`. KIND 2 is able to prove the satisfaction of the three properties, however KIND 2 reasoning is oblivious to the fact that there are two guarantees in the contract of `Display_Control` that make the specification unrealizable. The new feature of KIND 2 for checking the realizability of contracts is able to detect that.

In addition, KIND 2 provides a deadlocking computation and a set of conflicting guarantees to help the designer identify the source of the problem. In particular, in this case KIND 2 returns a deadlocking computation where initially all inputs are *false* and `minutes_to_cook` is 0, and, at the next step, both `cancel` and `decr` are *true*, the rest of inputs are *false*, and `minutes_to_cook` is 599. It also reports that, for the provided deadlocking computation, guarantees G5 and G9 form a minimal set of conflicting guarantees since, when both `cancel` and `decr` are *true* simultaneously, `minutes_to_cook` cannot always be zero (G5) and be decreased by one (G9) at the same time. One way to fix this issue is to update guarantee G9 to strengthen the premises of the implication with the requirement that the `cancel` button is not pressed. If the realizability of the contract is analyzed again after this update, an analogous conflict is found between G5 and G8. Figure 4

```

1 node Since ( X, Y : bool ) returns ( Z : bool ) ;
2 let
3   Z = X or (Y and (false -> pre Z)) ;
4 tel
5
6 node Once(x : bool) returns (Y : bool);
7 let
8   Y = x or (false -> pre Y);
9 tel
10
11 node Display_Control_Observer(
12   cancel: bool; incr: bool; decr: bool; baking: bool
13 )
14 returns(
15   left_digit: digit_range; middle_digit: digit_range; right_digit: digit_range;
16   minutes_to_cook: int
17 );
18 (*@contract
19   var ctime: int = (left_digit*60) + (middle_digit*10) + right_digit;
20
21   guarantee "P1: minutes_to_cook shall match the total time of the displayed digits"
22     minutes_to_cook = ctime;
23   guarantee "P2: minutes_to_cook shall not exceed the MAX_TIME value"
24     minutes_to_cook <= MAX_TIME;
25   guarantee "P3: If ctime is equal one, and since ctime was greater or equal than two, incr
26     has not been pressed and the oven has not been baking, then decr must have been pressed
27     once in the past"
28     Since(ctime>=2, not incr and not baking) and ctime=1 => Once(decr);
29 *)
30 let
31   left_digit, middle_digit, right_digit, minutes_to_cook =
32     Display_Control(cancel, incr, decr, baking);
33 tel

```

Fig. 3: KIND 2's contract for the Observer component

shows guarantees G8 and G9 after both contracts have been updated. The contract with the revised guarantees is proven realizable by KIND 2. Moreover, KIND 2 is still able to prove properties P1-P3.

2 Preliminaries

Lustre is a synchronous dataflow language that allows one to define system components as *nodes*, each of which maps a continuous stream of inputs (of various basic types, such as Booleans, integers, and reals) to continuous streams of outputs based on both current input values and previous input and output values. Bigger components can be built by parallel composition of smaller ones, achieved syntactically with *node applications*. Operationally, a node has a cyclic behavior: at each tick t of a global clock (or a local clock it is explicitly associated with) it reads the value of each input stream at position or *time* t , and instantaneously computes and returns the value of each output stream at time t .

Formally, a stream of values of type τ is a function the natural numbers (modeling the global clock ticks) to τ . The behavior of a Lustre node is then specified declaratively by a set of stream

```

1 guarantee "G8: When not baking, if incr is pressed but not cancel, minutes_to_cook shall
   increase by one if it was less than MAX or be zero otherwise"
2 true ->
3   (not baking and not cancel and incr) =>
4   (minutes_to_cook = if pre minutes_to_cook < MAX then pre minutes_to_cook + 1 else 0);
5
6 guarantee "G9: When not baking, if decr is pressed but not cancel nor incr, minutes_to_cook
   shall decrease by one if it was greater than 0 or be MAX otherwise"
7 true ->
8   (not baking and not cancel and not incr and decr) =>
9   (minutes_to_cook = if pre minutes_to_cook > 0 then pre minutes_to_cook - 1 else MAX);

```

Fig. 4: Guarantees G8 and G9 updated to make Display_Control’s contract realizable

constraints of the form $x = s$, where x is a variable denoting an output or a locally defined stream and s is a stream term over input, output, and local variables. Most stream operators are point-wise liftings of the usual operators over stream values. For example, if x and y are two integer streams, the term $x + y$ is the stream corresponding the function $\lambda t. x(t) + y(t)$ over time t ; an integer constant c , denotes the constant function $\lambda t. c$. Two important additional operators are a unary right-shift operator **pre**, used to specify stateful computations, and a binary initialization operator **->**, used to specify initial state values. At time $t = 0$, the value $(\text{pre } x)(t)$ is undefined; for each time $t > 0$, it is $x(t - 1)$. In contrast, the value $(x \text{ -> } y)(t)$ equals $x(t)$ for $t = 0$ and $y(t)$ for $t > 0$. Syntactic restrictions guarantee that all streams in a node are inductively well defined. In Kind 2’s extension of Lustre, nodes can be given assume-guarantee contracts, enabling the compositional analysis of Lustre models. Contracts specify assumptions as Boolean terms over current values of input streams and previous values of input and output streams, and guarantees as Boolean terms over current and previous values of input and output streams.

After various transformations and slicing, KIND 2 encodes Lustre nodes internally as (state) transition systems $S = \langle \mathbf{s}, \mathbf{i}, I[\mathbf{s}, \mathbf{i}], T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \rangle$ where \mathbf{s} is a vector of typed state variables, \mathbf{i} is a vector of typed input variables, I is the initial state predicate (over the variables in \mathbf{s} and \mathbf{i}), and T is a two-state transition predicate (over the variables in \mathbf{s} , \mathbf{i}' and \mathbf{s}' , with \mathbf{i}' and \mathbf{s}' being a renamed version of \mathbf{s} and \mathbf{i} , respectively). System outputs are represented by selected elements of \mathbf{s} which we do not distinguish from internal state variables for simplicity. We will use $\langle I, T \rangle$ to refer to transition system S when the vectors of state and input variables, \mathbf{s} and \mathbf{i} , are clear from the context or not important. A *contract* for S is a pair $C = \langle A, G \rangle$ of an *assumption* transition system $A = \langle \mathbf{i}, \mathbf{s}, A_I[\mathbf{i}], A_T[\mathbf{s}, \mathbf{i}'] \rangle$, where \mathbf{i} and \mathbf{s} act, respectively, as the state variables and the input variables of the environment of S ,¹ and a *guarantee* transition system $G = \langle \mathbf{s}, \mathbf{i}, G_I[\mathbf{s}, \mathbf{i}], G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \rangle$ with the same state and input variables as transition system S . The predicates A_I and A_T specify, respectively, which inputs are *valid* initially and for a given system state. The predicate G_I specifies which states the system may start in when the initial inputs satisfy A_I . The predicate G_T specifies for a given state and inputs what states the system may transition to when the inputs satisfy A_T . Given a contract $C = \langle A, G \rangle$, we will assume that G_I and G_T have the structure of a top-level conjunction, that is, $G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] = G_{T1}[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \wedge \dots \wedge G_{Tn}[\mathbf{s}, \mathbf{i}', \mathbf{s}']$ for some $n \geq 1$. Notice that Kind 2’s assume-guarantee contracts follow naturally this kind of conjunctive structure since they

¹ For simplicity, but without loss of generality, we assume that any reference to a previous value of an input variable in A_T is made through a state variable in \mathbf{s} storing that value, and thus, A_T is defined only over \mathbf{s} and \mathbf{i}' .

are specified as conjunction of assumptions and a conjunction of guarantees. By a slight abuse of notation, we will identify G_I (G_T), with the set $\{G_{I1}, \dots, G_{Im}\}$ ($\{G_{T1}, \dots, G_{Tn}\}$) of its top-level conjuncts.

Now we will introduce some definitions and results required to describe the new functionality of KIND 2. Given a vector of typed variables \mathbf{x} , a *valuation* ν over \mathbf{x} is a type-consistent assignment of values to all the variables in \mathbf{x} . For a valuation ν over a vector \mathbf{x} , we denote by ν' the valuation such that $\nu'(x') = \nu(x)$ for all variables x in \mathbf{x} ; for a vector \mathbf{y} consisting of variables from \mathbf{x} , we denote by $\nu[\mathbf{y}]$ the valuation over \mathbf{y} obtained by restricting ν to the variables in \mathbf{y} . Given a transition system $S = \langle \mathbf{s}, \mathbf{i}, I[\mathbf{s}, \mathbf{i}], T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \rangle$, a *state* of S is a valuation over \mathbf{s} and an *input* of S is a valuation over \mathbf{i} . A *trace* σ is a sequence of valuations over \mathbf{s} and \mathbf{i} . A *computation path* π of S of length $k \geq 0$ is a finite sequence of valuations $\pi = \pi_0, \dots, \pi_k$ over \mathbf{s} and \mathbf{i} such that π_0 satisfies the predicate $I[\mathbf{s}, \mathbf{i}]$, and for every $0 \leq i < k$ the valuations π_i, π'_{i+1} satisfy the predicate $T[\mathbf{s}, \mathbf{i}', \mathbf{s}']$. A *trace of S* is a trace $\sigma = \sigma_0, \sigma_1, \dots$ such that for every $k \geq 0$ the prefix $\sigma_0, \dots, \sigma_k$ is a computation path of S . We will denote the set of all the traces of S as $L(S)$. Given two transition systems S_1 and S_2 with the same vectors of state and inputs variables, S_2 is a (trace-based) *refinement* of S_1 iff $L(S_2) \subseteq L(S_1)$.

The following definitions are adapted from similar notions introduced by Gacek et al. [5]. Unlike the original paper, we explicitly formalize the fact that predicate G_I may depend on input values, and that assumptions may specify constraints over the initial input values through the predicate A_I . This allows for greater generality and flexibility without significantly affecting the context or the proven results. In the following, we fix for convenience a transition system $S = \langle \mathbf{s}, \mathbf{i}, I[\mathbf{s}, \mathbf{i}], T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \rangle$. In the definitions below, we will consider

- assumptions A of the form $\langle \mathbf{i}, \mathbf{s}, A_I[\mathbf{i}], A_T[\mathbf{s}, \mathbf{i}'] \rangle$ and
- guarantees G of the form $\langle \mathbf{s}, \mathbf{i}, G_I[\mathbf{s}, \mathbf{i}], G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \rangle$

for S .

Definition 1 A computation path $\pi = \pi_0, \dots, \pi_k$ of S satisfies an assumption A if π_0 satisfies the predicate $A_I[\mathbf{i}]$, and for every $0 \leq i < k$ the valuations π_i, π'_{i+1} satisfies the predicates $A_T[\mathbf{s}, \mathbf{i}']$.

Definition 2 A state $\hat{\mathbf{s}}$ is reachable (in S) under an assumption A if there exists a computation path $\pi = \pi_0, \dots, \pi_k$ of S satisfying A such that $\pi_k[\mathbf{s}] = \hat{\mathbf{s}}$. Formally, the set $\text{REACHABLE}_{S,A}(\mathbf{s})$ of reachable states under an assumption A is defined inductively by the following equation:

$$\text{REACHABLE}_{S,A}(\mathbf{s}) \triangleq (\exists \mathbf{i}_0. A_I[\mathbf{i}_0] \wedge I[\mathbf{s}]) \vee (\exists \mathbf{s}_p, \mathbf{i}'. \text{REACHABLE}_{S,A}(\mathbf{s}_p) \wedge A_T[\mathbf{s}_p, \mathbf{i}'] \wedge T[\mathbf{s}_p, \mathbf{i}', \mathbf{s}])$$

Definition 3 The transition system S satisfies a contract $C = \langle A, G \rangle$ when the following conditions hold:

1. $\forall \mathbf{s}, \mathbf{i}. A_I[\mathbf{i}] \wedge I[\mathbf{s}, \mathbf{i}] \Rightarrow G_I[\mathbf{s}, \mathbf{i}]$
2. $\forall \mathbf{s}, \mathbf{i}', \mathbf{s}'. \text{REACHABLE}_{S,A}(\mathbf{s}) \wedge A_T[\mathbf{s}, \mathbf{i}'] \wedge T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \Rightarrow G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}']$

When S does not satisfy a contract $C = \langle A, G \rangle$, there is a computation path $\pi = \pi_0, \dots, \pi_k$ of S satisfying assumption A such that either $k = 0$ and π_0 does not satisfy $G_I[\mathbf{s}, \mathbf{i}]$ or $k > 0$, π_0, \dots, π_{k-1} is a computation path of G but π_{k-1}, π'_k does not satisfy $G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}']$. We call such computation path a *safety counterexample*, and any trace that has that computation path as a prefix a *safety counter-trace*.

Definition 4 The transition system S is input-enabled under an assumption A when the following two conditions hold:

1. $\forall \mathbf{i}. A_I[\mathbf{i}] \Rightarrow \exists \mathbf{s}. I[\mathbf{s}, \mathbf{i}]$
2. $\forall \mathbf{s}, \mathbf{i}'. \text{REACHABLE}_{S,A}(\mathbf{s}) \wedge A_T[\mathbf{s}, \mathbf{i}'] \Rightarrow \exists \mathbf{s}'. T[\mathbf{s}, \mathbf{i}', \mathbf{s}']$

Definition 5 The transition system S is a realization of a contract $C = \langle A, G \rangle$ if S satisfies C and is input-enabled under assumption A .

Definition 6 A contract is realizable if there exists a transition system which is a realization of the contract.

When a contract $C = \langle A, G \rangle$ is unrealizable, we can try to build an *environment* transition system $E = \langle \mathbf{i}, \mathbf{s}, E_I[\mathbf{i}], E_T[\mathbf{s}, \mathbf{i}'] \rangle$ such that E is a realization of contract $\langle \langle \mathbf{s}, \mathbf{i}, \top, \top \rangle, A \rangle$, thus E is a refinement of A which always keeps running, and $L(E) \subseteq \overline{L(G)}$. We call E a *counter-strategy*. A user can examine a counter-strategy to try understand the reasons the contract is unrealizable and fix it accordingly. However, as pointed out by Könighofer et al [8], a counter-strategy may be very large and complex. Hence, the user may prefer a single computation path $\pi = \pi_0, \dots, \pi_k$ of G satisfying A such that state $\pi_k[\mathbf{s}]$ satisfies

$$\exists \mathbf{i}'. A_T[\mathbf{s}, \mathbf{i}'] \wedge \forall \mathbf{s}'. \neg G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \quad (1)$$

We say that $\pi_k[\mathbf{s}]$ is a *deadlocked* state.

Since knowing concrete input values for the existentially quantified variables in (1) is relevant to understand why G_T cannot be satisfied, instead of giving the user the computation path above, we return an extended version of it. Namely, our algorithm generates computation path $\hat{\pi} = \pi_0, \dots, \pi_k, \pi_{k+1}$ such that $\pi_k[\mathbf{s}]$ and $\pi'_{k+1}[\mathbf{i}']$ satisfy $A_T[\mathbf{s}, \mathbf{i}'] \wedge \forall \mathbf{s}'. \neg G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}']$. When an initial state does not always exist, the algorithm, instead, generates a computation path $\rho = \rho_0$ such that $\rho_0[\mathbf{i}']$ satisfies $A_I[\mathbf{i}'] \wedge \forall \mathbf{s}. \neg G_I[\mathbf{s}, \mathbf{i}']$. We will call such a computation path a *deadlocking computation*, and any trace that has the computation path as a prefix a *realizability counter-trace*. Although $\pi'_{k+1}(\rho_0)$ may give arbitrary values to $\mathbf{s}'(\mathbf{s})$, our algorithm computes it so that a minimal set U of guarantee conjuncts are violated, where $U \subseteq G_I$ when the violation happens at the initial step, and $U \subseteq G_T$ when it happens later. We call such a subset a *set of conflicting guarantees* or, simply, a *conflict*.

The realizability check presented in this paper is based on a notion called *viability* introduced by Gacek et al. [5], which provides a characterization of contract realizability.

Definition 7 A state $\hat{\mathbf{s}}$ is viable with respect to a contract C , if G_T can keep responding to valid inputs forever, starting from $\hat{\mathbf{s}}$. Formally, the set of viable states with respect to C is defined coinductively by the following equation:

$$\text{VIABLE}_C(\mathbf{s}) \triangleq \forall \mathbf{i}'. A_T[\mathbf{s}, \mathbf{i}'] \Rightarrow \exists \mathbf{s}'. G_T[\mathbf{s}, \mathbf{i}', \mathbf{s}'] \wedge \text{VIABLE}_C(\mathbf{s}')$$

Theorem 1. A contract C is realizable if and only if $\forall \mathbf{i}. A_I[\mathbf{i}] \Rightarrow \exists \mathbf{s}. G_I[\mathbf{s}, \mathbf{i}] \wedge \text{VIABLE}_C(\mathbf{s})$ holds.

Proof. Follows from our definition of input-enabled transition system and an analogous proof to the one provided for Theorem 1 in [5].

3 An Algorithm for Checking Realizability

In this section we present the algorithm used by KIND 2 for automatically checking the realizability of a contract, and finding a deadlocking computation and a conflict when the contract is proven unrealizable. It is an adaptation to KIND 2 of a synthesis procedure by Katis et al. [7]. The algorithm iteratively refines an over-approximation of the set of viable states, expressed as a predicate F , until F is determined to be a fixpoint by proving the validity of the following formula:

$$\forall \mathbf{s}, \mathbf{i}'. (F[\mathbf{s}] \wedge A_T[\mathbf{s}, \mathbf{i}'] \Rightarrow \exists \mathbf{s}'. G_T[\mathbf{s}, \mathbf{i}, \mathbf{s}'] \wedge F[\mathbf{s}']) \quad (2)$$

After the greatest fixpoint is computed, the realizability of the contract can be established by checking whether for all initial valid inputs there exists a state that satisfies G_I and F . When that is the case, the contract is realizable. Otherwise, the contract is unrealizable.

To decide the validity of $\forall\exists$ -formulas, the main algorithm relies on the AE-VAL procedure (described in Algorithm 1). AE-VAL starts computing a *region of validity* for the input formula, i.e., a formula $P[\mathbf{x}]$ such that $\forall \mathbf{x}. Q[\mathbf{x}] \wedge P[\mathbf{x}] \Rightarrow \exists \mathbf{y}. W[\mathbf{x}, \mathbf{y}]$ is valid. It achieves that by applying quantifier elimination to the formula $\exists \mathbf{y}. Q[\mathbf{x}] \Rightarrow W[\mathbf{x}, \mathbf{y}]$ which takes into account the context $Q[\mathbf{x}]$ (line 1). Then, it checks whether the formula $P[\mathbf{x}]$ is valid by checking if its negation is unsatisfiable. If it is, the original formula is valid. Otherwise, the original formula is invalid. In both cases, the algorithm conjoins the computed region of validity with $Q[\mathbf{x}]$ and then, it returns the region together with a Boolean value indicating the validity result.

Algorithm 1 AE-VAL ($\forall \mathbf{x}. Q[\mathbf{x}] \Rightarrow \exists \mathbf{y}. W[\mathbf{x}, \mathbf{y}]$)

- 1: $P[\mathbf{x}] \leftarrow \text{QE}(\exists \mathbf{y}. Q[\mathbf{x}] \Rightarrow W[\mathbf{x}, \mathbf{y}])$
 - 2: $\text{SmtAssert}(\neg P[\mathbf{x}])$
 - 3: $\text{sat} \leftarrow \text{SmtCheckSat}()$
 - 4: **return** $\langle \neg \text{sat}, Q[\mathbf{x}] \wedge P[\mathbf{x}] \rangle$
-

The realizability check procedure is described in Algorithm 2. It begins by checking that there exists a state satisfying G_I for all initial valid inputs (line 2). When that is not the case, the contract is unrealizable and a deadlocking computation is generated together with a set of conflicting guarantees (line 4). This check can be seen as an optimization for detecting unrealizable contracts without having to compute F , but it also helps to handle separately the generation of a deadlocking computation for the initial case and the transition case (line 19). Then, the algorithm checks whether the contract is trivially realizable because there are no initial valid inputs (line 6). If it is the case, the algorithm terminates declaring the contract realizable (line 7). Otherwise, it initializes four variables before entering the main loop (line 8): F represents the current candidate fixpoint, fl is a flag that indicates whether F has been refined at least once, and R and \hat{R} are used to store (after the first refinement) the regions of validity over \mathbf{s} and \mathbf{i}' , and \mathbf{s} , respectively, for which there exists a next state satisfying G_T . Both R and \hat{R} are arbitrarily initialized to \top .

In each iteration, the algorithm proceeds as follows. First, it checks whether greatest fixpoint has been reached by checking the validity of Formula 2 (line 11). If the formula is invalid, AE-VAL provides a region of validity $validRegion[\mathbf{s}, \mathbf{i}']$ over \mathbf{s} and \mathbf{i}' . This formula may contain constraints over the contract's inputs, so it cannot be used to refine F directly. To determine the specific region over \mathbf{s} for which there exists an input that violates Formula 2, we can use AE-VAL again

to determine the validity of formula $\phi' \leftarrow \forall \mathbf{s}. (F[\mathbf{s}] \Rightarrow \exists \mathbf{i}'. A_T[\mathbf{s}, \mathbf{i}'] \wedge \neg \text{validRegion}[\mathbf{s}, \mathbf{i}'])$. The invalidity of ϕ' indicates that there are still non-violating states (i.e., outside $\text{violatingRegion}[\mathbf{s}]$) which may lead to a fixpoint. Thus, the algorithm removes the unsafe states from $F[\mathbf{s}]$ in line 24, and iterates until a greatest fixpoint for $F[\mathbf{s}]$ is reached. If ϕ' is valid, then every state in $F[\mathbf{s}]$ is unsafe, under a specific input that satisfies the contract assumptions (since $\neg \text{validRegion}[\mathbf{s}, \mathbf{i}']$ holds in this case), and the specification is unrealizable. In the next iteration, the algorithm will reach line 20. In addition, when fl is *false* (i.e. it is the first iteration), the algorithm records $\text{validRegion}[\mathbf{s}, \mathbf{i}']$ and $\neg \text{violatingRegion}[\mathbf{s}]$ which are used to generate a deadlocking computation and a conflict if the contract is determined to be unrealizable.

If ϕ is valid, the algorithm checks whether for all initial valid inputs there exists a state that satisfies G_I and F (line 14). If so, the the contract is realizable and the algorithm returns the generated fixpoint (line 16). Otherwise, the contract is unrealizable and the algorithm generates a deadlocking computation and a set of conflicting guarantees. Since the algorithm has already verified that it is always possible to compute an initial state for any valid initial input (line 3), any counter-trace must involve one or more states. Moreover, F must has been refined at least once, and R and \hat{R} set to the regions of validity over \mathbf{s} and \mathbf{i}' , and \mathbf{s} , respectively, for which there exists a next state satisfying G_T (lines 26-27). To generate the deadlocking computation, we must find a computation path of G satisfying A that reaches a state $\hat{\mathbf{s}}$ from which it is impossible to transition to a new state satisfying G_T , i.e. $\models \neg \hat{R}[\hat{\mathbf{s}}]$. To find such computation path, the algorithm relies on a `Verify` procedure that receives a transition system $S = \langle I, T \rangle$ and a contract C , and returns a pair $\langle r, c \rangle$ where r indicates whether S satisfies the contract C or not, and c is a safety counterexample when S does not satisfy C . The algorithm use `Verify` to check whether transition system $\langle G_I, G_T \rangle$ satisfy contract $\langle A, \hat{R} \rangle$. Because the contract is unrealizable, it is ensured that the call to `Verify` in line 18 always determines that S does not satisfy C and it returns a counterexample satisfying the properties stated above.

To help the user to understand why a contract is unrealizable, Algorithm 3 computes a set of conflicting guarantees, and a valuation for the inputs and the state variables such that it satisfies as many guarantees as possible either initially, when the check in line 3 of Algorithm 2 was invalid, or from the final deadlocked state computed in line 18 of Algorithm 2 otherwise. This last valuation is appended to the deadlocking computation at the end.

Algorithm 3 first initializes ϕ with a constraint that defines the valuation of \mathbf{s} for the last state in the input counterexample, when the counterexample is not empty, or with \top otherwise. Then, it creates activation literals L (line 6) that will be used to track the contribution of each guarantee in G^* to the unsatisfiability of $A^* \wedge \varphi[\mathbf{s}] \wedge \neg R[\mathbf{s}, \mathbf{i}'] \wedge \bigwedge_{g_j \in G^*} l_j \Rightarrow g_j$. But first, the algorithm finds a valuation that maximizes the number of satisfied guarantees in G^* by solving a MaxSMT problem consisting in the hard constraint introduced in line 7, and a soft clause for each activation literal guarding a guarantee constraint (line 10). The algorithm uses the generated model θ to fix the values for the inputs in the last step (line 13). Then, it computes a minimal set of unsatisfiable guarantees (line 15). In lines 16-20 the algorithm extends the input counterexample with the computed valuation. Finally, the algorithm returns the final deadlocking computation and the set of conflicting guarantees based on the activation literals included in the unsat core (line 21).

4 Related Work

The realizability check described in this report is largely based on the synthesis procedure for infinite-state reactive systems, called JSYN-VG, presented in [7]. The only difference between both

Algorithm 2 REALIZABILITYCHECK ($A = \langle A_I, A_T \rangle, G = \langle G_I, G_T \rangle$)

```
1:  $\varphi \leftarrow \forall i. A_I[i] \Rightarrow \exists s. G_I[s, i]$ 
2:  $\langle valid, validRegion[i] \rangle \leftarrow \text{AE-VAL}(\varphi)$ 
3: if  $\neg valid$  then
4:    $cex, C \leftarrow \text{GETDEADLOCKINGCOMPANDCONFLICT}(validRegion[i], \text{EmptyList}(), A_I, G_I)$ 
5:   return  $\langle \text{UNREALIZABLE}, \langle cex, C \rangle \rangle$ 
6: if  $\text{IsUNSAT}(A_I[i])$  then
7:   return  $\langle \text{REALIZABLE}, \top \rangle$ 
8:  $F[s] \leftarrow \top; R[s, i'] \leftarrow \top; \hat{R}[s] \leftarrow \top; fl \leftarrow false$ 
9: while  $true$  do
10:   $\phi \leftarrow \forall s, i. (F[s] \wedge A_T[s, i'] \Rightarrow \exists s'. G_T[s, i, s'] \wedge F[s'])$ 
11:   $\langle valid, validRegion[s, i'] \rangle \leftarrow \text{AE-VAL}(\phi)$ 
12:  if  $valid$  then
13:     $\phi' \leftarrow \forall i. A_I[i] \Rightarrow \exists s. G_I[s, i] \wedge F[s]$ 
14:     $\langle valid', \_ \rangle \leftarrow \text{AE-VAL}(\phi')$ 
15:    if  $valid'$  then
16:      return  $\langle \text{REALIZABLE}, F[s] \rangle$ 
17:    else
18:       $\_, cex \leftarrow \text{Verify}(\langle G_I, G_T \rangle, \langle A, \hat{R} \rangle)$ 
19:       $cex', C \leftarrow \text{GETDEADLOCKINGCOMPANDCONFLICT}(R, cex, A_T, G_T)$ 
20:      return  $\langle \text{UNREALIZABLE}, \langle cex', C \rangle \rangle$ 
21:  else
22:     $\phi' \leftarrow \forall s. (F[s] \Rightarrow \exists i'. A_T[s, i'] \wedge \neg validRegion[s, i'])$ 
23:     $\langle \_, violatingRegion[s] \rangle \leftarrow \text{AE-VAL}(\phi')$ 
24:     $F[s] \leftarrow F[s] \wedge \neg violatingRegion[s]$ 
25:    if  $\neg fl$  then
26:       $R \leftarrow validRegion[s, i']$ 
27:       $\hat{R} \leftarrow \neg violatingRegion[s]$ 
28:       $fl \leftarrow false$ 
29: end while
```

Algorithm 3 GETDEADLOCKINGCOMPANDCONFLICT (R, cex, A^*, G^*)

```
1: if IsEmptyList( $cex$ ) then ▷ Initial Case
2:    $\varphi[\mathbf{s}] \leftarrow \top$ 
3: else ▷ Transition Case
4:    $\sigma \leftarrow \text{GetLastElementOfList}(cex)$  ▷ Map from state and input variables to values
5:    $\varphi[\mathbf{s}] \leftarrow \bigwedge_{s_j \in \mathbf{s}} s_j = \sigma(s_j)$ 
6: Create activation literals  $L = \{l_j \mid g_j \in G^*\}$ 
7: SmtAssert( $A^* \wedge \varphi[\mathbf{s}] \wedge \neg R[\mathbf{s}, \mathbf{i}'] \wedge \bigwedge_{g_j \in G^*} l_j \Rightarrow g_j$ )
8: SmtPush()
9: for  $l_j \in L$  do
10:   SmtAssertSoft( $l_j, 1$ )
11:  $\theta \leftarrow \text{SmtCheckSatAndGetModel}()$ 
12: SmtPop()
13: SmtAssert( $\bigwedge_{i_j \in \mathbf{i}} i_j = \theta(i_j)$ )
14:  $\_ \leftarrow \text{SmtCheckSatAssuming}(L)$ 
15:  $U \leftarrow \text{SmtGetMinimalUnsatCore}()$ 
16: if IsEmptyList( $cex$ ) then
17:    $\sigma' \leftarrow \{i_j \mapsto \theta(i_j) \mid i_j \in \mathbf{i}\} \cup \{s_j \mapsto \theta(s_j) \mid s_j \in \mathbf{s}\}$ 
18: else
19:    $\sigma' \leftarrow \{i_j \mapsto \theta(i_j) \mid i_j \in \mathbf{i}\} \cup \{s'_j \mapsto \theta(s'_j) \mid s'_j \in \mathbf{s}'\}$ 
20:  $cex' \leftarrow \text{AddElementAtTheEnd}(\sigma', cex)$ 
21: return  $cex', \text{MapActivationLiteralsToGuarantees}(U)$ 
```

works is more practical than theoretical. While the original work relies on a dedicated solver to implement the functionality provided by the AE-VAL procedure [4], our tool only requires a generic quantifier elimination procedure for the underlying theories supported by KIND 2 (LIA and LRA). These procedures are commonly available in state-of-the-art SMT solvers like Z3 [10] and CVC4 [1]. The use of a standard solver is also the approach followed by the synthesis tool GENSYS, recently published in [11], which was developed contemporary with our tool. As the experimental evaluation shows later, the use of standard solvers can improve the performance and increase the set of solved instances on the set of benchmarks used in the original work.

Another notable realizability check algorithm for infinite-state specifications is the one presented in [5], called JSYN, which follows a k-induction approach. Like the algorithm described in this report, it is also based on the notion of viability explained in Section 2. However, the algorithm suffers from soundness problems with respect to unrealizable results which limits its applicability.

A recent work on realizability checking of infinite-state specifications is the compositional realizability analysis presented in [9], which is a preprocessing step that can be applied to assume-guarantee contracts. It automatically partitions specifications into sets of non-interfering requirements so that checking whether a specification is realizable reduces to checking that each partition is realizable. Since this is an orthogonal technique that can improve the scalability of the functionality provided by KIND 2, we will study its integration in KIND 2 in the future.

5 Experimental Evaluation

We compared our realizability check implementation in KIND 2 with the latest version of JSYN-VG available within the JKind model checker (<https://github.com/andrewkatis/jkind-1/releases/tag/1.8>). We ran each tool on a Linux machine with eight 4-core Intel i7-6700 processors and 32GB

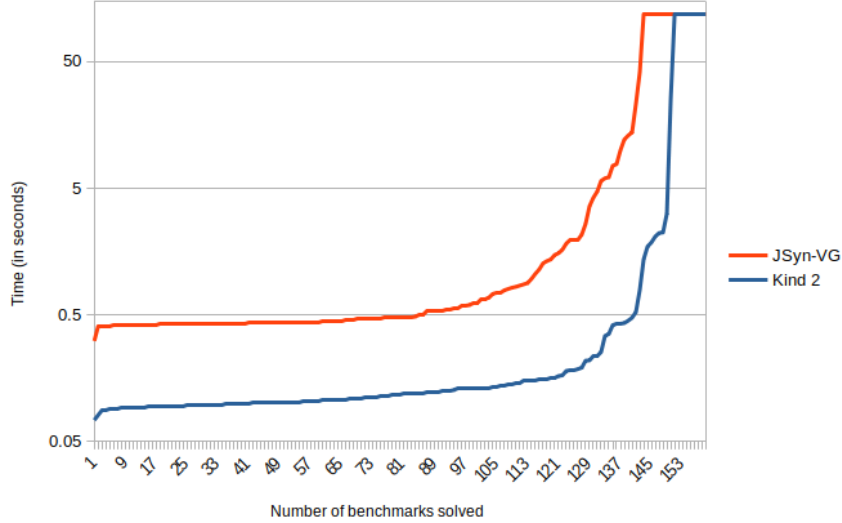


Fig. 5: Comparison between JKIND (JSYN-VG) and KIND 2

of memory using a timeout of 2 minutes. We used the benchmarks available at <https://github.com/andreaskatis/synthesis-benchmarks>, which includes the 124 contracts used in [7] plus 50 more contracts added to the repository after the publication of the work. Since JKIND doesn't have native support for the specification of contracts, the benchmarks are encoded using Lustre `assert` statements, and two special statements, `REALIZABLE` and `PROPERTY`. To run KIND 2 on the benchmarks, we encoded the problems using KIND 2 built-in assume-guarantee specification language.

After running the experiments, we found that KIND 2 rejected two of the problems before any analysis was performed due to syntactic restrictions imposed by KIND 2, and that KIND 2 and JKIND disagreed upon the result on 13 of the problems. The two rejected problems contained assumptions over current values of outputs streams, which KIND 2 does not accept as a way of encouraging good practices when writing specifications. We have often found that this kind of assumptions are not usually what the user intended to specify and they lead to subtle flaws. With regard to the problems where KIND 2 and JKIND disagreed on, they included unguarded applications of the `pre` operator, which leads to undefined behavior at the initial step. In the semantics of JKIND, each unguarded applications of the `pre` operator over the *same* expression is treated as a *single* undefined constant value. In contrast, KIND 2's semantics treats each unguarded applications of the `pre` operator as a potentially *different* undefined constant value even if it is applied to the same expression. This leads KIND 2 to classify as unrealizable problems that JKIND classifies as realizable.

To make a fair comparison we decided to remove the 15 problems mentioned above from the set of benchmarks, and carry our experimental evaluation over the remaining 159 problems. Moreover, the experimental evaluation only takes into account the runtime required to determine the realizability of the contracts, and thus, it excludes the generation of the deadlocking computation and conflict in the case of KIND 2, and the synthesis of an implementation in the case of JKIND.

Figure 5 shows that KIND 2 out-performances the implementation of JSYN-VG in JKIND providing an answer faster and in more cases. Moreover, the set of problems solved by KIND 2 is a

strictly larger superset of the problems solved by JKIND. When we doubled the original timeout up to 4 minutes, JKIND was able to solve only one more problem already solved by KIND 2.

In addition, we quantified the overhead of generating a deadlocking computation and a conflict for the contracts on the benchmark set that KIND 2 classified as unrealizable. Computing the additional information for the 22 contracts which KIND 2 could prove unrealizable took 37 seconds more, increasing the total runtime from 87 to 124 seconds. This represents a 43% overhead.

References

1. Barrett, C.W., Conway, C.L., Deters, M., Hadarean, L., Jovanovic, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Gopalakrishnan, G., Qadeer, S. (eds.) Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6806, pp. 171–177. Springer (2011). https://doi.org/10.1007/978-3-642-22110-1_14, https://doi.org/10.1007/978-3-642-22110-1_14
2. Champion, A., Gurfinkel, A., Kahsai, T., Tinelli, C.: Cocospec: A mode-aware contract language for reactive systems. In: Nicola, R.D., eua Kühn (eds.) Software Engineering and Formal Methods - 14th International Conference, SEFM 2016, Held as Part of STAF 2016, Vienna, Austria, July 4-8, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9763, pp. 347–366. Springer (2016). https://doi.org/10.1007/978-3-319-41591-8_24
3. Champion, A., Mebsout, A., Stickse, C., Tinelli, C.: The Kind 2 model checker. In: Chaudhuri, S., Farzan, A. (eds.) Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9780, pp. 510–517. Springer (2016). https://doi.org/10.1007/978-3-319-41540-6_29
4. Fedyukovich, G., Gurfinkel, A., Gupta, A.: Lazy but effective functional synthesis. In: Enea, C., Piskac, R. (eds.) Verification, Model Checking, and Abstract Interpretation - 20th International Conference, VMCAI 2019, Cascais, Portugal, January 13-15, 2019, Proceedings. Lecture Notes in Computer Science, vol. 11388, pp. 92–113. Springer (2019). https://doi.org/10.1007/978-3-030-11245-5_5, https://doi.org/10.1007/978-3-030-11245-5_5
5. Gacek, A., Katis, A., Whalen, M.W., Backes, J., Cofer, D.D.: Towards realizability checking of contracts using theories. In: Havelund, K., Holzmann, G.J., Joshi, R. (eds.) NASA Formal Methods - 7th International Symposium, NFM 2015, Pasadena, CA, USA, April 27-29, 2015, Proceedings. Lecture Notes in Computer Science, vol. 9058, pp. 173–187. Springer (2015). https://doi.org/10.1007/978-3-319-17524-9_13, https://doi.org/10.1007/978-3-319-17524-9_13
6. Halbwachs, N., Lagnier, F., Ratel, C.: Programming and verifying real-time systems by means of the synchronous data-flow language LUSTRE. IEEE Trans. Software Eng. **18**(9), 785–793 (1992). <https://doi.org/10.1109/32.159839>
7. Katis, A., Fedyukovich, G., Guo, H., Gacek, A., Backes, J., Gurfinkel, A., Whalen, M.W.: Validity-guided synthesis of reactive systems from assume-guarantee contracts. In: Beyer, D., Huisman, M. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part II. Lecture Notes in Computer Science, vol. 10806, pp. 176–193. Springer (2018). https://doi.org/10.1007/978-3-319-89963-3_10, https://doi.org/10.1007/978-3-319-89963-3_10
8. Könighofer, R., Hofferek, G., Bloem, R.: Debugging formal specifications: a practical approach using model-based diagnosis and counterstrategies. Int. J. Softw. Tools Technol. Transf. **15**(5-6), 563–583 (2013). <https://doi.org/10.1007/s10009-011-0221-y>, <https://doi.org/10.1007/s10009-011-0221-y>
9. Mavridou, A., Katis, A., Giannakopoulou, D., Kooi, D., Pressburger, T., Whalen, M.W.: From partial to global assume-guarantee contracts: Compositional realizability analysis in FRET. In: Huisman, M., Pasareanu, C.S., Zhan, N. (eds.) Formal Methods - 24th International Symposium, FM 2021, Virtual Event, November 20-26, 2021, Proceedings. Lecture Notes in Computer Science, vol. 13047, pp. 503–523. Springer (2021). https://doi.org/10.1007/978-3-030-90870-6_27, https://doi.org/10.1007/978-3-030-90870-6_27
10. de Moura, L.M., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4963, pp. 337–340. Springer (2008). https://doi.org/10.1007/978-3-540-78800-3_24, https://doi.org/10.1007/978-3-540-78800-3_24

11. Samuel, S., D'Souza, D., Komondoor, R.: Gensys: a scalable fixed-point engine for maximal controller synthesis over infinite state spaces. In: Spinellis, D., Gousios, G., Chechik, M., Penta, M.D. (eds.) ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23-28, 2021. pp. 1585–1589. ACM (2021). <https://doi.org/10.1145/3468264.3473126>, <https://doi.org/10.1145/3468264.3473126>