

A Theory of Synchronous Relational Interfaces

STAVROS TRIPAKIS and BEN LICKLY, University of California, Berkeley
THOMAS A. HENZINGER, Institute of Science and Technology, Austria
EDWARD A. LEE, University of California, Berkeley

Dedicated to the Memory of Amir Pnueli.

Compositional theories are crucial when designing large and complex systems from smaller components. In this work we propose such a theory for synchronous concurrent systems. Our approach follows so-called interface theories, which use game-theoretic interpretations of composition and refinement. These are appropriate for systems with distinct inputs and outputs, and explicit conditions on inputs that must be enforced during composition. Our interfaces model systems that execute in an infinite sequence of synchronous rounds. At each round, a contract must be satisfied. The contract is simply a relation specifying the set of valid input/output pairs. Interfaces can be composed by parallel, serial or feedback composition. A refinement relation between interfaces is defined, and shown to have two main properties: (1) it is preserved by composition, and (2) it is equivalent to substitutability, namely, the ability to replace an interface by another one in any context. Shared refinement and abstraction operators, corresponding to greatest lower and least upper bounds with respect to refinement, are also defined. Input-complete interfaces, that impose no restrictions on inputs, and deterministic interfaces, that produce a unique output for any legal input, are discussed as special cases, and an interesting duality between the two classes is exposed. A number of illustrative examples are provided, as well as algorithms to compute compositions, check refinement, and so on, for finite-state interfaces.

Categories and Subject Descriptors: D.2.2 [Software Engineering]: Design Tools and Techniques—*Modules and interfaces*; D.2.13 [Software Engineering]: Reusable Software

General Terms: Algorithms, Design, Languages, Theory, Verification

Additional Key Words and Phrases: Compositionality, interfaces, refinement, substitutability

ACM Reference Format:

Tripakis, S., Lickly, B., Henzinger, T. A., and Lee, E. A. 2011. A theory of synchronous relational interfaces. *ACM Trans. Program. Lang. Syst.*, 33, 4, Article 14 (July 2011), 41 pages.
DOI = 10.1145/1985342.1985345 <http://doi.acm.org/10.1145/1985342.1985345>

This report is a revised version of Tripakis et al. [2009a, 2009b].

This work was supported in part by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF awards #CCR-0225610 (ITR), #0720882 (CSR-EHS: PRET) and #0931843 (ActionWebs)), the U. S. Army Research Office (ARO #W911NF-07-2-0019), the U. S. Air Force Office of Scientific Research (MURI #FA9550-06-0312 and AF-TRUST #FA9550-06-1-0244), the Air Force Research Lab (AFRL), the Multiscale Systems Center (MuSyC) and the following companies: Bosch, National Instruments, Thales, and Toyota. This work was also supported by the COMBEST and ArtistDesign projects of the European Union, the Swiss National Science Foundation, the ERC Advanced Grant QUAREM and the FWF NFN Grant S11402-N23 (RiSE).

Author's address: S. Tripakis, 545Q, DOP Center, Cory Hall, EECS Department, University of California, Berkeley, CA 94720-1772; email: {stavros, blickly, eal}@eecs.berkeley.edu, tah@ist.ac.at.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permission may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2011 ACM 0164-0925/2011/07-ART14 \$10.00

DOI 10.1145/1985342.1985345 <http://doi.acm.org/10.1145/1985342.1985345>

1. INTRODUCTION

Compositional methods, which allow one to assemble smaller components into larger systems both efficiently and correctly, are not simply a desirable feature in system design: they are a must for designing large and complex systems. A compositional theory provides means for reasoning formally about components and their compositions. It also typically provides sufficient and/or necessary conditions for *substitutability*: when can a certain component be replaced by another one without compromising the correctness of the overall system? This property is clearly extremely important, in particular for incremental design.

The goal of this work is to develop a compositional theory for *synchronous concurrent systems*, systems where a set of components execute in an infinite sequence of *global rounds*. This is a fundamental model of computation with traditionally strong application in the hardware domain (digital circuits). Today the synchronous paradigm is also becoming widespread in software, in particular, in the domains of embedded and cyber-physical systems [Henzinger and Sifakis 2007; Lee 2008]. Tools such as Simulink from The MathWorks,¹ SCADE from Esterel Technologies,² or Ptolemy from Berkeley,³ and languages such as the synchronous languages [Benveniste et al. 2003] are important players in this field [Miller et al. 2010]. The semantics used in these models are synchronous.

Our work is situated in the context of *interface theories* [de Alfaro and Henzinger 2001a,b]. An interface can be seen as an abstraction of a component: on one hand, it captures information that is essential in order to use the component in a given context; on the other hand, it hides unnecessary information, making reasoning simpler and more efficient. Interface theories typically define a set of composition operators and a refinement relation on interfaces, and provide theorems of preservation of correctness by refinement and preservation of refinement by composition, from which substitutability guarantees can be derived. These concepts are common to most compositional theories. What distinguishes interface theories is a *game-theoretic* interpretation of the basic concepts, namely, composition and refinement. The need for a game-theoretic interpretation has been argued extensively in previous works on interface theories [de Alfaro 2004; de Alfaro and Henzinger 2001a, 2001b]. In order to make this article more self-contained, we also discuss the motivations behind this choice here, in Section 2.

The type of information about a component that is exposed in an interface varies depending on the application. For instance, in standard programming languages such as C or Java, the signature of a given method can be seen as an interface for that method. This interface provides sufficient information for type checking, but usually does not provide enough information for more detailed analysis, for instance, checking that a method computing division never attempts a division by zero. As this simple example illustrates, we should not expect a single “fits-all” interface theory, but multiple theories that are more or less suitable for different purposes. Suitability metrics could include expressiveness and ease of modeling in particular application domains, as well as tractability of the computational problems involved.

In our theory, an interface consists of a set of input variables X , a set of output variables Y , and a set of *contracts*. Semantically, a contract is simply a set of assignments of values to variables in $X \cup Y$. Syntactically, we use a logical formalism such as first-order logic to represent and manipulate contracts. For example, the predicate

¹<http://www.mathworks.com/products/simulink/>

²<http://www.esterel-technologies.com/products/scade-suite/>

³<http://ptolemy.eecs.berkeley.edu/>

$x_2 \neq 0 \wedge y = \frac{x_1}{x_2}$ can be used to represent the contract of a component that performs division, with input variables x_1 and x_2 and output variable y . The contract here is the set of all assignments to variables x_1, x_2 and y that satisfy the predicate. The assignment $(x_1 := 6, x_2 := 2, y := 3)$ satisfies the contract, while any assignment where $x_2 := 0$ violates the contract. A more abstract contract for the same component, which only gives some information about the sign of the output based on the sign of the inputs, is $x_2 \neq 0 \wedge (y < 0 \equiv (x_1 < 0 < x_2 \vee x_2 < 0 < x_1))$. An even more abstract contract is $x_2 \neq 0$. The latter guarantees nothing about the output; however, it still enforces that requirement that when performing division the denominator should be nonzero. We should note that these contracts implicitly use the fact that variables are numbers, symbols like $=$ for equality, and arithmetic operations such as division. Our theory does not depend on these, and it works with variables of any domain, without assuming any properties on such domains. In practice, however, we often use such properties implicitly for convenience.

Contracts govern the operation of a component, which evolves in a sequence of synchronous rounds. Within a round, values are assigned to the input variables of the component by its environment, and the component assigns values to its output variables. Together the two assignments form a complete assignment over all variables. This assignment must satisfy the contract. A new assignment is found at each round. Interfaces can be *stateless* or *stateful*. In the stateless case, there is a single contract that must hold at every round (the assignments may still differ). In the general, stateful case, a different contract may be specified for each round. The contract in this case depends on the history of assignments observed so far, which we call a *state*. The set of states, as well as the set of contracts, can be infinite. When the set of contracts is finite, we have a *finite-state interface* (note that the domains of variables could still be infinite). Finite-state interfaces are represented as finite automata whose locations are labeled by contracts.

Interfaces can be composed by *connection* or by *feedback* (see Section 6). Connection essentially corresponds to serial (cascade) composition, however, it can also capture parallel composition as a special case (empty connection). Composition by connection is generally not the same as composition of relations. Section 2 discusses this choice extensively. Feedback is allowed only for *Moore interfaces*, where the contract does not depend on the current values of the input variables that are back-fed (although it may depend on past values of such variables). A *hiding* operator (Section 7) can be used to eliminate redundant or intermediate output variables. Hiding is always possible for stateless interfaces and corresponds to existentially quantifying variables in the contract. The situation is more subtle in the stateful case, where we need to ensure that the “hidden” variables do not influence the evolution of the contract from one state to the next. This is necessary to ensure preservation of refinement by hiding.

Our theory includes explicit notions of *environments*, *pluggability*, and *substitutability* (Section 8). An environment E for an interface I is simply an interface whose input and output variables “mirror” those of I . I is pluggable to E (and vice versa) iff the closed-loop system formed by connecting the two is *well-formed*, that is, never reaches a state with an unsatisfiable contract. Substitutability means that an interface I' can replace another interface I in any environment. That is, for any environment E , if I is pluggable to E then I' is also pluggable to E .

Our refinement relation is similar in spirit to existing relations, such as function subtyping in type theory [Pierce 2002], behavioral subtyping [Liskov and Wing 1994], conformation in trace theory [Dill 1987], and alternating refinement [Alur et al. 1998]. All these, roughly speaking, state that I' refines I if I' accepts more

inputs and produces fewer outputs than I . This requirement is easy to formalize as $\text{in} \rightarrow \text{in}' \wedge \text{out}' \rightarrow \text{out}$ when the input assumptions, in , are separated from the output guarantees, out . When the constraints on the inputs and outputs are mixed in the same contract ϕ , a more careful definition is needed, namely: $\text{in}(\phi) \rightarrow (\text{in}(\phi') \wedge (\phi' \rightarrow \phi))$, where $\text{in}(\phi)$ characterizes the set of *legal* input assignments specified by ϕ . An input assignment is legal if there exists an output assignment such that together the two assignments satisfy the contract. For example, if ϕ is $x_2 \neq 0 \wedge y = \frac{x_1}{x_2}$ then $\text{in}(\phi)$ is $x_2 \neq 0$.

This definition of refinement applies to the stateless case where an interface has a single contract ϕ . The definition can be extended to the stateful case as shown in Section 9. Refinement is a partial order with the following main properties: (1) it is preserved by composition and hiding; and (2) it is essentially equivalent to substitutability (Theorem 9.14). It is worth noting that reasonable alternative definitions of refinement result in sufficient but not necessary conditions for substitutability (see discussions in Sections 2.4 and 9.2). Our notion of refinement thus is as strong as necessary for substitutability, but not stronger.

Our theory supports *shared refinement* (Section 10), which is important for component reuse as argued in Doyen et al. [2008]. Shared refinement of two interfaces I and I' , when defined, is a new interface that refines both I and I' , in fact, it is their greatest lower bound with respect to the refinement order, and is therefore denoted $I \sqcap I'$. In this article we also propose *shared abstraction* $I \sqcup I'$, which is shown to be the least upper bound with respect to refinement.

As a special case, we discuss *input-complete* (sometimes also called *receptive*) interfaces, where contracts are total relations, and *deterministic* interfaces, where contracts are partial functions. These two subclasses of interfaces are interesting, first, because the theory is greatly simplified in those cases (refinement becomes language containment, composition becomes relational, etc.), and second, because there is an interesting duality between the two subclasses, as shown in Sections 11 and 12.

Examples illustrating the concepts of the theory are provided throughout the article. An application to the hardware domain is described in Section 13.

The main features of the theory are summarized in Table I. This table is given merely for reference and contains only a partial view. The precise definitions and complete set of results are given in the sections that follow.

One of the appealing features of our theory is that it allows a *declarative* way of specifying contracts, and a *symbolic* way of manipulating them, as logical formulas. For this reason, it is relatively straightforward to develop algorithms that implement the theory for finite-state interfaces. Throughout the text we provide such algorithms, for composing interfaces, checking refinement, and so on. These algorithms compute some type of product of the automata that represent the interfaces and syntactically manipulate their contracts. Solving problems such as quantifier elimination and satisfiability checking on the formulas representing the contracts are crucial elements of the algorithms. Decidability of these problems will of course depend on the types of formulas used. Recent advances in SMT (Satisfiability Modulo Theory) solvers can be leveraged for this task.

2. MOTIVATION FOR THE DESIGN CHOICES

As mentioned in the introduction, our theory uses a game-theoretic interpretation of composition and refinement. These interpretations are by no means new and have been motivated in previous works (see Section 3). In this section we also motivate these choices in our setting.

Table I. Summary of the Main Concepts and Results of This Article

| | | |
|---------------------------------|---|---|
| Models | Relational interfaces | Stateless: contracts (predicates) on input/output variables, e.g., $x_2 \neq 0 \wedge y = \frac{x_1}{x_2}$. Legal input assignments in contract ϕ : $\text{in}(\phi) := (\exists y_1, y_2, \dots, y_n : \phi)$, where $\{y_1, \dots, y_n\}$ is the set of output variables, e.g., $\text{in}(x_2 \neq 0 \wedge y = \frac{x_1}{x_2}) \equiv x_2 \neq 0$. |
| | | Stateful: automata whose states are labeled with contracts (set of states may be infinite). |
| | Moore w.r.t. input x | Contract does not depend on current value of variable x (can still depend on previous values of x). |
| | Input-complete | All input values are legal: $\text{in}(\phi) \equiv \text{true}$. |
| | Deterministic | Given legal inputs, outputs are unique. |
| | Well-formed | All reachable contracts are satisfiable. |
| | Well-formable | Can be made well-formed by restricting the inputs. |
| Environments | They are interfaces. | |
| Compositions | Connection | Parallel: conjunction of contracts: $\phi := \phi_1 \wedge \phi_2$. Commutative & associative. Serial: game, environment vs. source interface: $\phi := \phi_1 \wedge \phi_2 \wedge \forall y_1, \dots, y_n : (\phi_1 \rightarrow \text{in}(\phi_2))$, where $\{y_1, \dots, y_n\}$ is the set of output variables of the source interface. Associative. |
| | Feedback | Interface must be Moore w.r.t. input variable x that is connected to output y . Commutative & associative. |
| | Pluggability: $I \rightleftharpoons E$ | Closed-loop composition of I and E must be well-formed |
| | Substitutability: $I \rightarrow_e I'$ | For any E , if I is pluggable to E then I' is pluggable to E |
| Compositionality | Refinement | For stateless: $\phi' \sqsubseteq \phi := (\text{in}(\phi) \rightarrow (\text{in}(\phi') \wedge (\phi' \rightarrow \phi)))$. Similar for stateful. \sqsubseteq is partial order. false is top element. |
| | Preservation | Refinement preserves well-formability. |
| | | Refinement is preserved by both connection and feedback. |
| | | Refinement sufficient for substitutability: if $I' \sqsubseteq I$ then $I \rightarrow_e I'$. |
| | | Refinement necessary for substitutability, in that: if $I' \not\sqsubseteq I$ and I is well-formed, then $I \not\rightarrow_e I'$. |
| Special case: input-complete | If ϕ_2 is input-complete then serial composition is conjunction. Refinement becomes: $\phi' \sqsubseteq \phi \equiv (\phi' \rightarrow \phi)$. | |
| Special case: deterministic | If ϕ_1 is deterministic then serial composition is conjunction. Refinement becomes: $\phi' \sqsubseteq \phi \equiv (\phi \rightarrow \phi')$. | |

2.1 A General Model for Contracts: Relational, Nondeterministic, Non-Input-Complete

Consider a component performing division. One possible interface for this component is the following:

$$Div := (\{x_1, x_2\}, \{y\}, \phi_{Div}^1)$$

$$\phi_{Div}^1 := x_2 \neq 0 \wedge \phi_{sign}$$

$$\phi_{sign} := (y = 0 \equiv x_1 = 0) \wedge (y < 0 \equiv (x_1 < 0 < x_2 \vee x_2 < 0 < x_1)).$$

Div has two input variables x_1, x_2 , one output variable y , and a contract represented by the predicate ϕ_{Div}^1 . Interpreting x_1 to be the dividend and x_2 the divisor, and y to be

the result of the division, ϕ_{Div}^1 states that the divisor must be nonzero and also provides guarantees on the sign of the output depending on the sign of the inputs.

The following points are worth making about contract ϕ_{Div}^1 . First, it is relational, in the sense that the value of the output depends on the values of the inputs. A non-relational contract that could be used is, for instance, $x_2 \neq 0$, which represents only an assumption on the input. Another nonrelational contract, for a slightly more restrictive component that does not accept negative inputs, would be $x_1 \geq 0 \wedge x_2 > 0 \wedge y \geq 0$. This is nonrelational in the sense that the guarantee on the output does not depend on the value of the inputs. The second point about ϕ_{Div}^1 is that it is nondeterministic: the output y is not uniquely determined for a given input (unless $x_1 = 0$). The final point about ϕ_{Div}^1 is that it is non-input-complete: all inputs where $x_2 = 0$ are illegal in the sense that they violate the contract.

As this example illustrates, “rich” contracts, that is, relational, nondeterministic, and non-input-complete, arise even in simple situations. The need to capture relations between inputs and outputs should be clear: if we separate input assumptions from output guarantees (as done in Doyen et al. [2008], for instance) then we cannot state input-output properties about our system. The need for nondeterminism should also be clear: nondeterminism is useful when abstracting low-level details that would be too difficult to obtain or too expensive to use. For instance, in our example, we could use a deterministic contract for *Div*:

$$\phi_{Div}^2 := x_2 \neq 0 \wedge x_1 = y \cdot x_2.$$

But we may opt for ϕ_{Div}^1 , since ϕ_{Div}^1 can be handled by an SMT solver that can only deal with linear constraints, whereas ϕ_{Div}^2 cannot.


The need for non-input-completeness may be less obvious. Why can’t we replace the non-input-complete contract ϕ_{Div}^1 by the input-complete contract

$$\phi_{Div}^3 := x_2 \neq 0 \rightarrow \phi_{sign} \quad ?$$

There are several reasons. First, note that ϕ_{Div}^3 allows inputs where $x_2 = 0$ (since an implication $A \rightarrow B$ is trivially satisfied when A is false) and in that case the output y may take any value. However, the implicit assumption is that y *will* take *some* value. In other words, a “real” component (in SW or HW) that implements ϕ_{Div}^3 must be “input-complete” in the sense that it always produces some output, even when given illegal inputs. But not all real systems have this property. For example, a program may not terminate on illegal inputs; and a circuit may burn up if an incorrect voltage is applied to it. These systems are not “input-complete,” thus cannot be described by input-complete interfaces.

But even when a system is “input-complete,” we may still want to capture it with a non-input-complete interface. Indeed, suppose we connect *Div* to another component *C*, as shown to the left of Figure 1. Our intention is for *C* to output the constant 2, so that the composition implements a division by 2. But suppose that due to a design error *C* outputs zero instead. That is, the contract of *C* is $x_2 = 0$. Combining the latter with ϕ_{Div}^1 , by taking the conjunction of the two, gives false, which means that two interfaces are incompatible. Catching this incompatibility early on, that is, when attempting to compose *C* with *Div*, is useful, since it permits to localize and correct the error easily.

Suppose we used the input-complete contract ϕ_{Div}^3 , instead of ϕ_{Div}^1 . Then, the composition of *C* and *Div* would result in the contract true (after hiding variable x_2). How should we interpret this result? We cannot in general interpret true as indicating incompatibility, since it might simply be the result of lack of information, that is, trivial contracts. In a large system, there will generally be many components for which we

Fig. 1. Connecting C to Div (left) and D to Div (right).

have no information, and others for which we do. We want a systematic (or even automatic) method that distinguishes between “no information” and “incompatible composition.”

We could of course perform a “local” verification task, in order to prove that the composition of C and Div implements the intended division by 2, namely, the property $\phi_P := (y \cdot 2 = x_1)$. Contract true fails to imply this, which indicates an error. The problem with this approach is that it requires ϕ_P to be specified. This may not always be an easy task. First, formal verification may not be part of the design process. Second, even if it is, it may be the case that only a global, end-to-end specification is available, for the top-level component within which the composition of C and Div is embedded. “Decomposing” such a global specification into local specifications such as ϕ_P is not always straightforward.

With non-input-complete interfaces, such local specifications are not required. Instead, a compatibility check is performed to ensure that a composition such as the one between C and Div is valid. This is a more lightweight verification process, akin to type checking (but with types that are richer than usual).

In fact, the goal may not be verification at all, but rather *synthesis* of component interfaces, in a bottom-up fashion: given interfaces for atomic components C and Div , compute an interface for their composition. Instead of type-checking, this is akin to type-inference. Once an interface for a complete hierarchical model is computed, and assuming no incompatibilities have been found during the process, that interface can be checked against a global specification, if the latter is available. But the interface provides useful information that can be helpful even in the absence of such a specification.

2.2 On the Definition of Serial Composition

Consider again interface Div , and suppose we connect it to another interface D , as shown to the right of Figure 1. Suppose that the contract of D is $\phi_D := \text{true}$. If D abstracts a certain component, this may mean that we have no knowledge about this component (e.g., our automated abstraction tool gave us a trivial answer).

What should the contract of the composition of D and Div be? Standard composition of relations corresponds to taking the conjunction $\phi_D \wedge \phi_{Div}^1$, and then eliminating x_2 . This yields the formula $\exists x_2 : \phi_{Div}^1$, which is equivalent to the predicate $y = 0 \equiv x_1 = 0$, asserting that y is zero iff x_1 is zero. This assertion is satisfiable (there are values for y and x_1 that make it true), therefore, it would appear that the composition of D and Div is valid.

Now, suppose that we want to replace D by E , which has the same structure as D (i.e., same input and output variables) but a different contract $\phi_E := x_2 = 0$. ϕ_E provides stronger output guarantees than ϕ_D , and in any standard framework this means that E refines D (this is also true in our framework). But clearly the composition of E and Div is not valid.

This means that even though E refines D , we cannot ensure that E can replace D : indeed, even though we accepted the composition of D and Div as valid, the composition of E and Div is not valid. This violates one of the main properties of any

compositional theory, namely substitutability, which states that a component should be replaceable by any component that refines it.

We are therefore forced to revise our assumption that the composition of D and Div is valid. The problem is that we interpreted the non-determinism of D as “angelic,” or “controllable.” We should instead interpret it as “demonic,” or “uncontrollable.” We should accept the composition as valid only if there exist input values at x_3 for which it can be guaranteed that any possible output of D satisfies $x_2 \neq 0$. Since D does not guarantee anything, no such input at x_3 can be found. Therefore, the composition of D and Div should be considered invalid. Logically, we achieve this by adding the term $\forall x_2 : \phi_D \rightarrow x_2 \neq 0$ to the conjunction $\phi_D \wedge \phi_{Div}^1$, in the definition of the composite contract. The above term reduces to $\forall x_2 : \text{true} \rightarrow x_2 \neq 0$, or $\forall x_2 : x_2 \neq 0$, which is false.

2.3 On the Definition of Refinement

Once we accept the “demonic” interpretation of nondeterminism in composition, as described above, the choice of refinement appears to be inevitable. Indeed, we seek a refinement relation that is equivalent to substitutability. This means that refinement must be both sufficient for substitutability (i.e., if interface I' refines interface I , then I' can replace I in any context) as well as necessary (i.e., if I' does not refine I , then there is a context where I works but I' does not). As it turns out, the definition that has these properties is the following: contract ϕ' refines contract ϕ , denoted $\phi' \sqsubseteq \phi$, iff the condition $\text{in}(\phi) \rightarrow (\text{in}(\phi') \wedge (\phi' \rightarrow \phi))$ is satisfied for any input/output assignment (this is the simplified definition for stateless interfaces, the general definition is given in Section 9). Alternative definitions can be given, which result in sufficient but not necessary conditions for substitutability, as discussed in Sections 2.4 and 9.2.

Refinement is not the same as logical implication. As an example, consider three possible contracts for our division component:

$$\begin{aligned}\phi_{Div}^1 &:= x_2 \neq 0 \wedge (y = 0 \equiv x_1 = 0) \wedge (y < 0 \equiv (x_1 < 0 < x_2 \vee x_2 < 0 < x_1)) \\ \phi_{Div}^2 &:= x_2 \neq 0 \wedge x_1 = y \cdot x_2 \\ \phi_{Div}^4 &:= x_2 \neq 0 \rightarrow x_1 = y \cdot x_2.\end{aligned}$$

It can be verified that $\phi_{Div}^4 \sqsubseteq \phi_{Div}^2 \sqsubseteq \phi_{Div}^1$. Yet observe that $\phi_{Div}^2 \rightarrow \phi_{Div}^1$ and $\phi_{Div}^2 \rightarrow \phi_{Div}^4$.

2.4 Error-Complete Interfaces

Illegal inputs can also be captured using input-complete interfaces that have a special Boolean output variable e , denoting an “error”. Let ϕ be a contract over input and output variables $X \cup Y$. Let e be a new output variable, not in Y . The *error-completion* of ϕ can be defined as the input-complete contract $\text{EC}(\phi)$ over $X \cup Y \cup \{e\}$, which sets e to false when the input is legal for ϕ , and to true otherwise:

$$\text{EC}(\phi) := (\phi \wedge \neg e) \vee (\neg \text{in}(\phi) \wedge e). \quad (1)$$

For example, the error-completion of the division interface Div^1 yields:

$$\begin{aligned}Div_e^1 &:= (\{x_1, x_2\}, \{y, e\}, \phi_e), \text{ where } \phi_e := \text{EC}(\phi_{Div}^1) \\ \phi_e &\equiv (x_2 \neq 0 \wedge \phi_1 \wedge \neg e) \vee (x_2 = 0 \wedge e).\end{aligned}$$

We can retrieve ϕ from $\text{EC}(\phi)$ using the inverse transformation:

$$\text{EC}^{-1}(\phi_e) := (\exists e : \phi_e) \wedge (\forall Y \cup \{e\} : \phi_e \rightarrow \neg e). \quad (2)$$

It can be shown that, if ϕ_e is of the form described in (1), then $(\exists e : \phi_e) \equiv (\phi \vee \neg \text{in}(\phi))$ and $(\forall Y \cup \{e\} : \phi_e \rightarrow \neg e) \equiv \text{in}(\phi)$. That is, the first term is the *input-completion* of ϕ which adds all its illegal inputs in its domain, whereas the second term isolates the

legal inputs. The conjunction of these two terms gives ϕ . Therefore, for any contract ϕ , we have:

$$\phi \equiv \text{EC}^{-1}(\text{EC}(\phi)). \quad (3)$$

On the other hand, for contracts ϕ_e over $X \cup Y \cup \{e\}$, $\text{EC}(\text{EC}^{-1}(\phi_e))$ is not always equivalent to ϕ_e . For example, if $\phi_e := y \equiv e$ and y is an output, then $\text{EC}^{-1}(\phi_e) \equiv \text{false}$, and $\text{EC}(\text{EC}^{-1}(\phi_e)) \equiv e$. Indeed, EC is injective but not surjective. It is unclear what is the meaning of a contract such as $y \equiv e$. This contract appears to “misuse” the error variable e which is supposed to capture validity of inputs.

It appears that game-theoretic serial composition of two contracts can be performed as a sequence of steps: error-completion, standard relational composition, and inverse error-completion. We illustrate this with an example. Consider the composition of interfaces D and Div discussed above. D is already input-complete, so its error-completion is unnecessary ($\text{EC}(\phi_D)$ would still extend ϕ_D with an additional error output variable, but we omit this for the sake of simplicity). Let $\phi_e := \text{EC}(\phi_{Div}^1)$. Let ϕ_{err} be the relational composition of ϕ_D and ϕ_e , that is: $\phi_{err} := \exists x_2 : (\text{true} \wedge \phi_e)$. It can be verified that $\phi_{err} \equiv \text{true}$. Then:

$$\text{EC}^{-1}(\phi_{err}) \equiv (\exists e : \text{true}) \wedge (\forall e : \text{true} \rightarrow \neg e) \equiv \text{true} \wedge \text{false} \equiv \text{false}.$$

This is indeed the same as the result obtained in Section 2.2 and indicating that the composition of D and Div is invalid.

Refinement between two contracts is different from logical implication of their error-transformed versions, that is, $\phi_2 \sqsubseteq \phi_1$ is not equivalent to $\text{EC}(\phi_2) \rightarrow \text{EC}(\phi_1)$. In particular, although validity of $\text{EC}(\phi_2) \rightarrow \text{EC}(\phi_1)$ is a sufficient condition for $\phi_2 \sqsubseteq \phi_1$, it is not a necessary condition. To see why, consider two interfaces that only model assumptions on an input variable x , and having contracts $x > 0$ and true , respectively. true accepts more inputs than $x > 0$, therefore we have $\text{true} \sqsubseteq x > 0$. Now consider $\psi_1 := \text{EC}(x > 0)$ and $\psi_2 := \text{EC}(\text{true})$. We have:

$$\begin{aligned} \psi_1 &\equiv (x > 0 \wedge \neg e) \vee (x \leq 0 \wedge e) \\ \psi_2 &\equiv \neg e. \end{aligned}$$

Clearly, $\psi_2 \not\sqsubseteq \psi_1$. Because of Theorem 9.14, which states equivalence of refinement and substitutability, this example also shows that $\text{EC}(\phi_2) \rightarrow \text{EC}(\phi_1)$ is a sufficient but not necessary condition for substitutability.

In summary, it appears that: (1) Error-complete interfaces can be used to capture the same information as that contained in relational interfaces. However, the class of error-complete interfaces is larger, and some of these interfaces have no direct meaning as relational interfaces. Thus, relational interfaces appear to be a more “canonical” representation. Moreover, relational interfaces avoid the overhead of designating special error outputs whose semantics differ from other outputs. (2) Composition of relational interfaces can be defined as standard relational composition of their error-complete counterparts. However, in order to check whether such a composition is valid, the inverse transformation needs to be computed. This inverse transformation involves solving a game, therefore, the game-theoretic interpretation of composition is not avoided. (3) Implication of their error-complete counterparts is a strictly stronger condition than refinement/substitutability between two interfaces.

Based on these observations, it appears that our theory could be formulated essentially equivalently in terms of error-complete interfaces. We do not pursue this option, however, as relational interfaces without special error outputs seem more elegant to us. On the other hand, error-complete interfaces are worth studying in greater depth, since error variables can be used for additional purposes than simply indicating illegal

inputs. For instance, they may be used to indicate faulty behavior of a component. An in-depth study of these possibilities is beyond the scope of the current article and part of future work.

3. RELATED WORK

Most of the ideas upon which this work is based, such as stepwise refinement, interfaces, design-by-contract, and game semantics, are by no means new. The main contribution of this article is the application of these ideas to the development of a working theory for synchronous concurrent systems.

In particular, abstracting components in some mathematical framework that offers stepwise refinement and compositionality guarantees is an idea that goes back to the work of Floyd and Hoare on proving program correctness using pre- and post-conditions [Floyd 1967; Hoare 1969] and the work of Dijkstra and Wirth on stepwise refinement as a method for gradually developing programs from their specifications [Dijkstra 1972; Wirth 1971]. A pair of pre- and postconditions can be seen as a contract for a piece of sequential code. These ideas were further developed in a large number of works, including the Z notation [Spivey 1989], the B method [Abrial 1996], CLU [Liskov 1979], Eiffel and the design-by-contract paradigm [Meyer 1992], the refinement calculus [Back and Wright 1998], Larch [Cheon and Leavens 1994; Guttag and Horning 1993; Leavens 1994], and JML [Leavens and Cheon 2006]. Viewing programs as predicates or relations is also not new; for instance, see Hoare [1985], Parnas [1983], Frappier et al. [1998], and Kahl [2003].

The above works are primarily about sequential programs and therefore are not directly comparable with our framework which is about synchronous concurrent systems. For instance, our model has distinct notions of input and output variables, whereas sequential programs operate on a set of shared variables, that is, a global, shared state. A program can be modeled as a relation between values of these global variables before and after program execution, that is, “pre” and “post” variables. However, it seems that our composition operators cannot be directly mapped to those aiming to capture typical constructs in sequential programs, such as “if-then-else” or “while” statements. Consider feedback composition, for example. One might attempt to map this to some form of while statement. But while statements operate on the same set of global variables, which could be seen as a special case of feedback where there is a 1-1 correspondence between inputs and outputs (i.e., pre and post variables). In the general case, an arbitrary output variable can be connected to an arbitrary input, which seems to make denotational approaches such as lifting to powersets inapplicable.

In fact, many of the previous works start with a programming language in mind, and then define the pre/postconditions, abstractions, or interfaces, for this particular language. As is characteristically stated in Hoare [1985], programs are predicates, but not all predicates are programs. In contrast, our framework is “implementation-agnostic” in the sense that we are not concerned with whether components are implemented in HW or in SW, or in which programming language. For this reason, as well as the fact that nonimplementable predicates such as false may arise as a result of composition, we do not attempt to restrict the set of predicates that we consider as contracts.

One concern that naturally arises in sequential programs that contain “while” loops is program termination. How can nontermination be modeled when programs are captured by relations? This question has received a lot of attention in the literature (an excellent survey can be found in Nelson [1989]) and has also generated some controversy [Hehner and Parnas 1985]. Our take on this is simple: if a component S may not terminate on a given input value a , then the contract for S should reject a as illegal.

That is, the input-output relation for S is partial. In terms of Nelson’s classification, our model can be seen as an instance of the “partial correctness model” [Nelson 1989]. This model does not distinguish a component S that never terminates on input a , from another component S' that may or may not terminate on a . We do not worry about this loss of expressiveness, however, because in our context, all components must be guaranteed to terminate and produce a value at *every* synchronous round. As a result, if an input may result in nontermination, it is appropriate to consider this input as illegal. Therefore, S' can be safely abstracted by the same interface as S .

Despite these differences, our approach follows many of the principles advocated in the works mentioned above. In particular, we abide to the design-by-contract paradigm and the well-known principle of refinement by weakening the precondition and strengthening the postcondition (although “strengthening the post condition” must be defined carefully in the general, non-input-complete case, as discussed in Section 9.2). Also, we use a “demonic” interpretation of nondeterminism during composition, as some of the works above also do [Back and Wright 1998; Frappier et al. 1998]. Computing composition then amounts to finding strategies in a game, or equivalently, solving a controller synthesis problem [de Alfaro 2004].

Interfaces can be viewed as “rich,” behavioral types, as suggested in Lee and Xiong [2001] and de Alfaro and Henzinger [2001a]. Behavioral types have been studied in a number of works in the context of sequential and object-oriented programming [Dhara and Leavens 1996; Liskov and Wing 1994; Nierstrasz 1993]. Behavioral subtyping notions defined in the above works follow the same principle of input-contravariance/output-covariance also in our refinement, but there are subtle differences in their definitions. For instance, both the “Post-condition rule” $m_\sigma.post \Rightarrow m_\tau.post$ and the “Constraint rule” $C_\sigma \Rightarrow C_\tau$, defined in Figure 4 of Liskov and Wing [1994] as requirements for type σ to be a subtype of type τ , appear to follow the rule $\phi' \rightarrow \phi$ rather than the rule $(in(\phi) \wedge \phi') \rightarrow \phi$ which is used in our refinement. As explained in Section 9.2, $\phi' \rightarrow \phi$ is too strong in the sense that it is not a necessary condition for substitutability. Dhara and Leavens [1996] weaken the subtyping requirements of Liskov and Wing [1994], but maintain the $\phi' \rightarrow \phi$ rule.

The works mentioned so far focus on sequential programs. In a concurrency setting, a powerful compositional framework is FOCUS [Broy 1997; Broy and Stølen 2001]. FOCUS is a relational framework where specifications are relations on input-output *streams*. The FOCUS framework is in many respects more general than ours, in that it can capture relations that do not preserve the length of input streams. For this reason, FOCUS is applicable also to asynchronous systems. On the other hand, FOCUS targets mainly the input-complete case. I/O automata [Lynch and Tuttle 1989] are also related to our work, but are input-complete by definition. Reactive modules [Alur and Henzinger 1999] are also input-complete.

Dill’s *trace theory* is another compositional framework for concurrent systems, focusing on asynchronous concurrency and motivated in particular by the design of asynchronous circuits [Dill 1987]. In trace theory, a component is described using a pair of sets of traces, called *successes* and *failures*, for legal and illegal behaviors, respectively. A *trace* is a sequence of events, and an event is a change in the value of an input or output variable. Because no synchrony is assumed, the number of input and output events in a trace can be arbitrary. Trace theory considers *prefix-closed trace structures*, where the success and failure sets are prefix-closed regular sets, aimed at verification of safety properties, as well as *complete trace structures*, where these are general sets of infinite traces, aimed at liveness properties. Our theory is currently restricted to prefix-closed sets and therefore cannot handle liveness properties. However, it is worth noting that, contrary to prefix-closed trace structures, our theory avoids the problem of trivial implementations that achieve the specification by “doing nothing.”

In trace theory, refinement is called *conformation* and is achieved by restricting the set of failures as well as the global set of traces (failures can be turned into successes during refinement). Conformation follows the “accept more inputs, produce less outputs” principle that has later been studied in the context of alternating refinement relations [Alur et al. 1998]. It is worth noting that conformation induces a lattice on prefix-closed trace structures, whereas our refinement relation is only a partial order and in particular has no “bottom” element.

Like trace structures, the framework of interface automata [de Alfaro and Henzinger 2001a] uses an asynchronous model of concurrency. Compared to trace structures, interface automata are more “syntactic” in nature since the interface is the automaton itself (as opposed to, say, a set of traces that can be represented by an automaton). Modal interfaces [Raclet et al. 2010] also focus on asynchronous systems and work directly with an automata representation. It is an interesting question to what extent these automata-based models can be used to capture synchronous systems and input-output relations within a synchronous round. If possible to do so, the result would most likely have an operational flavor, contrary to our framework, which is of a more declarative, denotational and symbolic nature. For instance, to express variables with infinite domains in the above formalisms, one would typically need an infinite set of events; to express a relation such as $y = x + 1$ one would need an infinite set of transitions; and so on.

Our theory can be used as a behavioral type theory for Simulink and related models, in the spirit of Roy and Shankar [2010]. In the latter work, Simulink blocks are annotated with constraints on input and output variables much like the stateless contracts considered in our work. Our framework provides an extension of such types to the stateful case, as well as the formalization of compositions and refinement that are not considered in Roy and Shankar [2010].

Within the domain of interface theories, de Alfaro and Henzinger [2001b] define *relational nets*, which are networks of processes that nondeterministically relate input values to output values. [de Alfaro and Henzinger 2001b] does not provide an interface theory for the complete class of relational nets. Instead it provides interface theories for subclasses, in particular: *rectangular nets* which have no input-output dependencies; *total nets* which can have input-output dependencies but are input-complete; and *total and rectangular nets* which combine both restrictions above. The interfaces provided in de Alfaro and Henzinger [2001b] for rectangular nets are called *assume/guarantee (A/G) interfaces*. A/G interfaces form a strict subclass of the relational interfaces that we consider in this article: A/G interfaces separate the assumptions on the inputs from the guarantees on the outputs, and as such cannot capture input-output relations; on the other hand, every A/G contract can be trivially captured as a relational contract by taking the conjunction of the assume and guarantee parts. [de Alfaro and Henzinger 2001b] studies *stateless* A/G interfaces, while Doyen et al. [2008] study also *stateful* A/G interfaces, in a synchronous setting similar to the one considered in this article. [Doyen et al. 2008] also discusses *extended interfaces* which are essentially the same as the relational interfaces that we study in this article. However, difficulties with synchronous feedback loops (see discussion that follows) lead Doyen et al. [2008] to conclude that extended interfaces are not appropriate.

Chakrabarti et al. [2002] consider synchronous *Moore interfaces*, defined by two formulas ϕ_i and ϕ_o that specify the legal values of the input and output variables, respectively, at the *next* round, given the current state. This formulation does not allow description of relations between inputs and outputs within the same round, as our relational theory allows.

Both de Alfaro and Henzinger [2001b] and Doyen et al. [2008] can handle very general compositions of interfaces, that can be obtained via parallel composition and

arbitrary connection (similar to the denotational composition framework of Lee and Sangiovanni-Vincentelli [1998]). This allows, in particular, arbitrary feedback loops to be created. In a relational framework, however, synchronous feedback loops can be problematic, as discussed in Example 9.11 (see also Section 14).

Interface theories are naturally related to work on compositional verification, where the main purpose is to break down the task of checking correctness of a large model into smaller tasks, that are more amenable to automation. A very large body of research exists on this topic. Some of this work is based on an asynchronous, interleaving based concurrency model [Jonsson 1994; Misra and Chandy 1981; Stark 1985] some on a synchronous model [Grumberg and Long 1994; McMillan 1997], while others are done within a temporal logic framework [Abadi and Lamport 1995; Barringer et al. 1984]. Many of these works are based on the assume-guarantee paradigm, and they typically use some type of trace inclusion or simulation as refinement relation [Henzinger et al. 1998; Jones 1983; Shankar 1998; Stark 1985].

4. PRELIMINARIES, NOTATION

We use first-order logic (FOL) notation throughout the article. For an introduction to FOL, see, for instance, Tournakis [2008]. We use true and false for logical constants true and false, \neg , \wedge , \vee , \rightarrow , \equiv for logical negation, conjunction, disjunction, implication, and equivalence, and \exists and \forall for existential and universal quantification, respectively. We use $:=$ when defining concepts or introducing new notation: for instance, $x_0 := \max\{1, 2, 3\}$ defines x_0 to be the maximum of the set $\{1, 2, 3\}$.

Let V be a finite set of variables. A *property over V* is a FOL formula ϕ such that any free variable of ϕ is in V . The set of all properties over V is denoted $\mathcal{F}(V)$. Let ϕ be a property over V and V' be a finite subset of V , $V' = \{v_1, v_2, \dots, v_n\}$. Then, $\exists V' : \phi$ is shorthand for $\exists v_1 : \exists v_2 : \dots : \exists v_n : \phi$. Similarly, $\forall V' : \phi$ is shorthand for $\forall v_1 : \forall v_2 : \dots : \forall v_n : \phi$.

We will implicitly assume that variables are *typed*, meaning that every variable is associated with a certain *domain*. An *assignment over a set of variables V* is a (total) function mapping every variable in V to a certain value in the domain of that variable. The set of all assignments over V is denoted $\mathcal{A}(V)$. If a is an assignment over V_1 and b is an assignment over V_2 , and V_1, V_2 are disjoint, we use (a, b) to denote the combined assignment over $V_1 \cup V_2$. A formula ϕ is *satisfiable* iff there exists an assignment a over the free variables of ϕ such that a satisfies ϕ , denoted $a \models \phi$. A formula ϕ is *valid* iff it is satisfied by every assignment.

There is a natural mapping from formulas to sets of assignments, that is, from $\mathcal{F}(V)$ to $2^{\mathcal{A}(V)}$. In particular, a formula $\phi \in \mathcal{F}(V)$ can be interpreted as the set of all assignments over V that satisfy ϕ . Conversely, we can map a subset of $\mathcal{A}(V)$ to a formula over V , provided this subset is representable in FOL. Because of this correspondence, we use set-theoretic or logical notation, as is more convenient. For instance, if ϕ and ϕ' are formulas or sets of assignments, we write $\phi \wedge \phi'$ or $\phi \cap \phi'$ interchangeably.

If S is a set, S^* denotes the set of all finite sequences of elements of S . S^* includes the empty sequence, denoted ε . If $s, s' \in S^*$, then $s \cdot s'$ is the concatenation of s and s' . $|s|$ denotes the *length* of $s \in S^*$, with $|\varepsilon| = 0$ and $|s \cdot a| = |s| + 1$, for $a \in S$. If $s = a_1 a_2 \dots a_n$, then the i -th element of the sequence, a_i , is denoted s_i , for $i = 1, \dots, n$. A *prefix* of $s \in S^*$ is a sequence $s' \in S^*$ such that there exists $s'' \in S^*$ such that $s = s' \cdot s''$. We write $s' \leq s$ if s' is a prefix of s . $s' < s$ means $s' \leq s$ and $s' \neq s$. A subset $L \subseteq S^*$ is *prefix-closed* if for all $s \in L$, for all $s' \leq s$, $s' \in L$.

5. RELATIONAL INTERFACES

Definition 5.1 (Relational interface). A *relational interface* (or simply *interface*) is a tuple $I = (X, Y, f)$ where X and Y are two finite and disjoint sets of *input* and *output variables*, respectively, and f is a nonempty, prefix-closed subset of $\mathcal{A}(X \cup Y)^*$.

Note that $\mathcal{A}(X \cup Y)$ can be infinite. In the case variables in X and Y have finite domains, $\mathcal{A}(X \cup Y)$ is finite and can be seen as a finite alphabet. In that case, f is a nonempty, prefix-closed language over that alphabet.

We write $\text{InVars}(I)$ for X and $f(I)$ for f . We allow X or Y to be empty: if X is empty then I is a *source* interface; if Y is empty then I is a *sink*. An element of $\mathcal{A}(X \cup Y)^*$ is called a *state*. That is, we identify states with observation histories. The *initial state* is the empty sequence ε . The states in f are also called the *reachable states* of I . f defines a total function that maps a state to a set of input-output assignments. We use the same symbol f to refer to this function. For $s \in \mathcal{A}(X \cup Y)^*$, $f(s)$ is defined as follows.

$$f(s) := \{a \in \mathcal{A}(X \cup Y) \mid s \cdot a \in f\}.$$

We view $f(s)$ as a *contract* between a component and its environment *at that state*. The contract changes dynamically, as the state evolves.

Conversely, if we are given a function $f : \mathcal{A}(X \cup Y)^* \rightarrow 2^{\mathcal{A}(X \cup Y)}$, we can define a non-empty, prefix-closed subset of $\mathcal{A}(X \cup Y)^*$ as follows:

$$f := \{a_1 \cdots a_k \mid \forall i = 1, \dots, k : a_i \in f(a_1 \cdots a_{i-1})\}$$

Notice that $\varepsilon \in f$ because the condition above trivially holds for $k = 0$. Also note that if $s \notin f$ then $f(s) = \emptyset$. This is because f is prefix-closed.

Because of the given 1-1 correspondence, in the sequel, we treat f either as a subset of $\mathcal{A}(X \cup Y)^*$ or as a function that maps states to contracts, depending on what is more convenient. We will assume that $f(s)$ is representable by a FOL formula. Therefore, $f(s)$ can be seen also as an element of $\mathcal{F}(X \cup Y)$.

Definition 5.2 (Input assumptions). Given a contract $\phi \in \mathcal{F}(X \cup Y)$, the *input assumption* of ϕ is the formula $\text{in}(\phi) := \exists Y : \phi$. Note that $\text{in}(\phi)$ is a property over X . Also note that $\phi \rightarrow \text{in}(\phi)$ is a valid formula for any ϕ .

A relational interface $I = (X, Y, f)$ can be seen as specifying a game between a component and its environment. The game proceeds in a sequence of *rounds*. At each round, an assignment $a \in \mathcal{A}(X \cup Y)$ is chosen, and the game moves to the next round. Therefore, the history of the game is the sequence of rounds played so far, that is, a state $s \in \mathcal{A}(X \cup Y)^*$. Suppose that at the beginning of a round the state is s . Typically, the environment plays first, by choosing $a_X \in \mathcal{A}(X)$. If $a_X \notin \text{in}(f(s))$ then this is not a legal input and the environment loses the game. Otherwise, the component plays by choosing $a_Y \in \mathcal{A}(Y)$. If $(a_X, a_Y) \notin f(s)$ then this is not a legal output for this input, and the component loses the game. Otherwise, the round is complete, and the game moves to the next round, with new state $s \cdot (a_X, a_Y)$. There are cases when the interface is *Moore* in the sense that its current outputs do not depend on its current inputs (the formal definition is given in Section 6.2). In this case, the component plays first. More general games can also be considered where the assignments of values to input and output variables are interleaved in an arbitrary order. The study of such a generalization is beyond the scope of the current work.

An *input-complete* interface is one that does not restrict its inputs.

Definition 5.3 (Input-complete interface). An interface $I = (X, Y, f)$ is *input-complete* if for all $s \in \mathcal{A}(X \cup Y)^*$, $\text{in}(f(s))$ is valid.

It is important to note that in our framework, input assumptions (“preconditions”) and output guarantees (“postconditions”) are not separated. It is then crucial to distinguish a non-input-complete interface with a contract of the form $\phi_{pre} \wedge \phi$ and its input-complete version with contract $\phi_{pre} \rightarrow \phi$. These contracts are different (as we will show in Section 11, the latter refines the former). As mentioned in Section 3, our theory is mostly implementation-agnostic, and therefore does not prescribe how illegal

inputs should be interpreted in the “real” component that an interface abstracts. As stated in Section 2, an illegal input may correspond to an input that results in non-termination of a SW component, or it may be an input that must be avoided by design, as in a type-checking setting.

A *deterministic* interface is one that maps every input assignment to at most one output assignment.

Definition 5.4 (Determinism). An interface $I = (X, Y, f)$ is *deterministic* if for all $s \in f$, for all $a_X \in \text{in}(f(s))$, there is a unique $a_Y \in \mathcal{A}(Y)$ such that $(a_X, a_Y) \in f(s)$.

The specializations of our theory to input-complete and deterministic interfaces are discussed in Sections 11 and 12, respectively.

A *stateless* interface is one where the contract is independent from the state.

Definition 5.5 (Stateless interface). An interface $I = (X, Y, f)$ is *stateless* if for all $s, s' \in \mathcal{A}(X \cup Y)^*$, $f(s) = f(s')$.

For a stateless interface, we can treat f as a subset of $\mathcal{A}(X \cup Y)$ instead of a subset of $\mathcal{A}(X \cup Y)^*$. For clarity, if I is stateless, we write $I = (X, Y, \phi)$, where ϕ is a property over $X \cup Y$.

Example 5.6 (Stateless interfaces). Consider a component which is supposed to take as input a positive number n and return n or $n + 1$ as output. We can capture such a component in different ways. One way is to use the following stateless interface:

$$I_1 := (\{x\}, \{y\}, x > 0 \wedge (y = x \vee y = x + 1)).$$

Here, x is the input variable and y is the output variable. The contract of I_1 explicitly forbids zero or negative values for x . Indeed, we have $\text{in}(f(I_1)) \equiv x > 0$.

Another possible stateless interface for this component is:

$$I_2 := (\{x\}, \{y\}, x > 0 \rightarrow (y = x \vee y = x + 1)).$$

The contract of I_2 is different from that of I_1 : it allows $x \leq 0$, but makes no guarantees about the output y in that case. I_2 is input-complete, whereas I_1 is not. Both I_1 and I_2 are nondeterministic.

In general, the state space of an interface is infinite. In some cases, however, only a finite set of states is needed to specify f . In particular, f may be specified by a finite-state automaton.

Definition 5.7 (Finite-state interface). A *finite-state interface* is specified by a finite-state automaton $M = (X, Y, L, \ell_0, C, T)$. X and Y are sets of input and output variables, respectively. L is a finite set of *locations* and $\ell_0 \in L$ is the initial location. $C : L \rightarrow 2^{\mathcal{A}(X \cup Y)}$ is a *labeling function* that labels every location with a set of assignments over $X \cup Y$, the contract at that location. $T \subseteq L \times 2^{\mathcal{A}(X \cup Y)} \times L$ is a set of *transitions*. A transition $t \in T$ is a tuple $t = (\ell, g, \ell')$ where ℓ, ℓ' are the source and destination locations, respectively, and $g \subseteq \mathcal{A}(X \cup Y)$ is the *guard* of the transition. We require that, for all $\ell \in L$:

$$C(\ell) = \bigcup_{(\ell, g, \ell') \in T} g \quad (4)$$

$$\forall (\ell, g_1, \ell_1), (\ell, g_2, \ell_2) \in T : \ell_1 \neq \ell_2 \rightarrow g_1 \cap g_2 = \emptyset. \quad (5)$$

These conditions ensure that there is a unique outgoing transition for every assignment that satisfies the contract of the location. Given $a \in C(\ell)$, the a -successor of ℓ is the unique location ℓ' for which there exists transition (ℓ, g, ℓ') such that $a \in g$. A

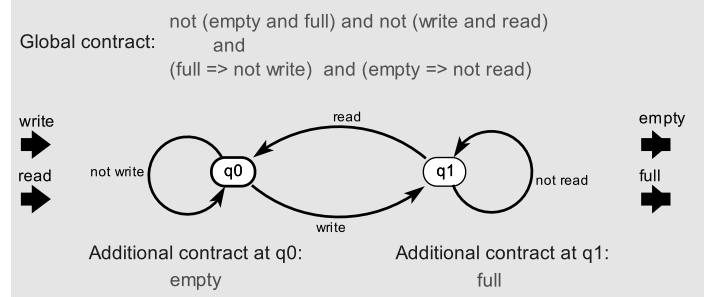


Fig. 2. Stateless and finite-state interfaces for a buffer of size 1.

location ℓ is called *reachable* if, either $\ell = \ell_0$, or there exists a reachable location ℓ' , a transition (ℓ', g, ℓ) , and an assignment a such that ℓ is the a -successor of ℓ' .

M defines interface $I = (X, Y, f)$ where f is the set of all sequences $a_1 \cdots a_k \in \mathcal{A}(X \cup Y)^*$, such that for all $i = 1, \dots, k$, $a_i \in C(\ell_{i-1})$, where ℓ_i is the a_i -successor of ℓ_{i-1} .


Note that a finite-state interface can still have variables with infinite domains. If the domains of variables are finite, however, then a finite-state interface can be seen as a prefix-closed regular language. Also notice that we allow $C(\ell)$, the contract at location ℓ , to be empty. This simply means that the interface is not well-formed (see Definition 5.9 that follows). Finally, although the guard of an outgoing transition from a certain location must be a subset of the contract of that location, we will often abuse notation and violate this constraint in the examples that follow, for the sake of simplicity. Implicitly, all guards should be understood in conjunction with the contracts of their source locations.

It is also worth noting that although the finite-state automaton defining a finite-state interface is deterministic, this does not mean that the interface itself is deterministic. Indeed, in general, it is not, since contracts that label locations are still nondeterministic input-output relations.

Example 5.8 (Stateful interface). Figure 2 shows a finite-state automaton defining a finite-state interface that captures a single-place buffer. The interface has two input variables, *write* and *read*, and two output variables, *empty* and *full*. All variables are boolean. The automaton has two locations, q_0 (the initial location) and q_1 . Each location is implicitly annotated by the conjunction of a *global contract*, that must hold at all locations, and a *local contract*, specific to a location. The global contract specifies that the buffer cannot be both empty and full (this is a guarantee on the outputs) and that a user of the buffer cannot read and write at the same round (this is an assumption on the inputs). The global contract also specifies that if the buffer is full then writing is not allowed, and if the buffer is empty then read is not allowed. The local contract at q_0 states that the buffer is empty and at q_1 that it is full.

Definition 5.9 (Well-formedness). An interface $I = (X, Y, f)$ is *well-formed* iff for all $s \in f$, $f(s)$ is nonempty.

Well-formed interfaces can be seen as describing components that never “deadlock.” If I is well-formed then for all $s \in f$ there exists assignment a such that $s \cdot a \in f$. Moreover, f is nonempty and prefix-closed by definition, therefore, $\varepsilon \in f$. This means that there exists at least one state in f which can be extended to arbitrary length. In a finite-state interface, checking well-formedness amounts to checking that the contract of every reachable location of the corresponding automaton is satisfiable. If contracts

Fig. 3. A well-formable interface I and its well-formed witness I' .

are specified in a decidable logic, checking well-formedness of finite-state interfaces is thus decidable.

Example 5.10. Let I be the finite-state interface represented by the leftmost automaton shown in Figure 3. I is assumed to have two Boolean variables, an input x , and an output y . I is not well-formed, because it has reachable states with contract false (all states starting with x being false). I can be transformed into a well-formed interface by strengthening the contract of the initial state from true to x , thus obtaining interface I' shown to the right of the figure.

Example 5.10 shows that some interfaces, even though they are not well-formed, can be turned into well-formed interfaces by appropriately restricting their inputs. This motivates the following definition.

Definition 5.11 (Well-formability). An interface $I = (X, Y, f)$ is *well-formable* if there exists a well-formed interface $I' = (X, Y, f')$ such that: for all $s \in f'$, $f'(s) \equiv f(s) \wedge \phi_s$, where ϕ_s is some property over X .

LEMMA 5.12. *Let $I = (X, Y, f)$ be a well-formable interface and let $I' = (X, Y, f')$ be a witness to the well-formability of I . Then $f' \subseteq f$.*

Proofs can be found in the online Appendix.

Clearly, every well-formed interface is well-formable, but the opposite is not true in general, as Example 5.10 shows. For stateless or source interfaces, however, the two notions coincide.

THEOREM 5.13. *A stateless or source interface I is well-formed iff it is well-formable.*

For an interface that is finite-state and whose contracts are written in a logic for which satisfiability is decidable, there is an algorithm to check whether the interface is well-formable, and if this is the case, to transform it into a well-formed interface. The algorithm essentially attempts to find a winning strategy in a *game*, and as such is similar in spirit to algorithms proposed in de Alfaro and Henzinger [2001a]. The algorithm starts by marking all locations with unsatisfiable contracts as *illegal*. Then, a location ℓ is chosen such that ℓ is legal, but has an outgoing transition (ℓ, g, ℓ') , such that ℓ' is illegal. If no such ℓ exists, the algorithm stops. Otherwise, the contract of ℓ is strengthened to

$$C(\ell) := C(\ell) \wedge (\forall Y : C(\ell) \rightarrow \neg g). \quad (6)$$

$\forall Y : C(\ell) \rightarrow \neg g$ is a property on X . An input assignment a_X satisfies this formula iff, for any possible output assignment a_Y that the contract $C(\ell)$ can produce given a_X , the complete assignment (a_X, a_Y) violates g . This means that there is a way of restricting the inputs at ℓ , so that ℓ' becomes unreachable from ℓ . Notice that, in the special case where g is a formula over X , (6) simplifies to $C(\ell) := C(\ell) \wedge \neg g$.

If, during the strengthening process, the contract of a location becomes unsatisfiable, this location is marked as illegal. The process is repeated until no more

strengthening is possible, whereupon the algorithm stops. Termination is guaranteed because each location has a finite number of successor locations, therefore, can only be strengthened a finite number of times. If, when the algorithm stops, the initial location ℓ_0 has been marked illegal, then the interface is not well-formed. Otherwise, the modified automaton specifies a well-formed interface, which is a witness for the original interface.

For this class of interfaces there is also an algorithm to check equality, that is, given two interfaces I_1, I_2 , check whether $I_1 = I_2$. Let $M_i = (X, Y, L_i, \ell_{0,i}, C_i, T_i)$ be finite-state automata representing I_i , for $i = 1, 2$, respectively. We first build a synchronous product $M := (X, Y, L_1 \times L_2 \cup \{\ell_{bad}\}, (\ell_{0,1}, \ell_{0,2}), C, T)$, where $C(\ell_1, \ell_2) := C_1(\ell_1) \vee C_2(\ell_2)$ for all $(\ell_1, \ell_2) \in L_1 \times L_2$, $C(\ell_{bad}) := \text{false}$, and:

$$T := \{((\ell_1, \ell_2), (C_1(\ell_1) \equiv C_2(\ell_2)) \wedge g_1 \wedge g_2, (\ell'_1, \ell'_2)) \mid (\ell_i, g_i, \ell'_i) \in T_i, \text{ for } i = 1, 2\} \\ \cup \{((\ell_1, \ell_2), C_1(\ell_1) \neq C_2(\ell_2), \ell_{bad})\} \quad (7)$$

It can be checked that $I_1 = I_2$ iff location ℓ_{bad} is unreachable.

6. COMPOSITION

We define two types of composition: by *connection* and by *feedback*.

6.1 Composition by Connection

First, we can compose two interfaces I_1 and I_2 “in sequence,” by connecting some of the output variables of I_1 to some of the input variables of I_2 . One output can be connected to many inputs, but an input can be connected to at most one output. Parallel composition is a special case of composition by connection, where the connection is empty. The connections define a new interface. Thus, the composition process can be repeated to yield arbitrary (for the moment, acyclic) interface diagrams. Composition by connection is associative (Theorem 6.6), so the order in which interfaces are composed does not matter.

Two interfaces $I = (X, Y, f)$ and $I' = (X', Y', f')$ are called *disjoint* if they have disjoint sets of input and output variables: $(X \cup Y) \cap (X' \cup Y') = \emptyset$.

Definition 6.1 (Composition by connection). Let $I_i = (X_i, Y_i, f_i)$, for $i = 1, 2$, be two disjoint interfaces. A *connection* θ between I_1, I_2 , is a finite set of pairs of variables, $\theta = \{(y_i, x_i) \mid i = 1, \dots, m\}$, such that: (1) $\forall (y, x) \in \theta : y \in Y_1 \wedge x \in X_2$, and (2) there do not exist $(y, x), (y', x) \in \theta$ such that y and y' are distinct. Define:

$$\text{InVars}(\theta) := \{x \mid \exists y : (y, x) \in \theta\} \quad (8)$$

$$X_{\theta(I_1, I_2)} := (X_1 \cup X_2) \setminus \text{InVars}(\theta) \quad (9)$$

$$Y_{\theta(I_1, I_2)} := Y_1 \cup Y_2 \cup \text{InVars}(\theta). \quad (10)$$

The connection θ defines the *composite interface* $\theta(I_1, I_2) := (X_{\theta(I_1, I_2)}, Y_{\theta(I_1, I_2)}, f)$, where, for every $s \in \mathcal{A}(X_{\theta(I_1, I_2)} \cup Y_{\theta(I_1, I_2)})^*$:

$$f(s) := f_1(s_1) \wedge f_2(s_2) \wedge \rho_\theta \wedge \forall Y_{\theta(I_1, I_2)} : \Phi \\ \Phi := (f_1(s_1) \wedge \rho_\theta) \rightarrow \text{in}(f_2(s_2)) \quad (11) \\ \rho_\theta := \bigwedge_{(y, x) \in \theta} y = x$$

and, for $i = 1, 2$, s_i is defined to be the projection of s to variables in $X_i \cup Y_i$.

Note that $X_{\theta(I_1, I_2)} \cup Y_{\theta(I_1, I_2)} = X_1 \cup Y_1 \cup X_2 \cup Y_2$. Also notice that $\text{InVars}(\theta) \subseteq X_2$. This implies that $X_1 \subseteq X_{\theta(I_1, I_2)}$, that is, every input variable of I_1 is also an input variable of $\theta(I_1, I_2)$. Also note that $\forall Y_{\theta(I_1, I_2)} : \Phi$ is equivalent to $\forall Y_1 \cup \text{InVars}(\theta) : \Phi$ because Φ

does not contain any Y_2 variables. The term $\forall Y_{\theta(I_1, I_2)} : \Phi$ is a condition on $X_{\theta(I_1, I_2)}$, the free inputs of the composite interface. This term states that, no matter which outputs I_1 chooses to produce for a given input, all such outputs are legal inputs for I_2 . This condition is essential for preservation of compatibility by refinement as discussed in Section 2.2, and more generally, for preservation of refinement by composition (Theorem 9.8).

Example 6.2. We repeat the example in Section 2.2 while being more pedantic. Let Div be the interface defined in Section 2.1 and let D be the interface $D := (\{x_3\}, \{y_2\}, \text{true})$. Let $\theta := \{(y_2, x_2)\}$. Then the term $\forall Y_{\theta(D, Div)} : \Phi$ instantiates to $\forall y_2, y : (\text{true} \wedge y_2 = x_2) \rightarrow x_2 \neq 0$, or equivalently, $\forall y_2 : y_2 \neq 0$, which is false, meaning that D and Div are “incompatible.” This notion is formalized next.

Contrary to other works [de Alfaro and Henzinger 2001a,b; Doyen et al. 2008], we do not impose an a-priori compatibility condition on connections. Not doing so allows us to state more general results (Theorem 9.8). Having said that, compatibility is a useful concept; therefore, we define it explicitly.

Definition 6.3 (Compatibility). Let I_1, I_2 be two disjoint interfaces and θ a connection between them. I_1, I_2 are said to be *compatible with respect to θ* iff $\theta(I_1, I_2)$ is well-formable.

For finite-state interfaces, connection is computable. Let $M_i = (X_i, Y_i, L_i, \ell_{0,i}, C_i, T_i)$ be finite-state automata representing I_i , for $i = 1, 2$, respectively. The composite interface $\theta(I_1, I_2)$ can be represented as $M := (X_{\theta(I_1, I_2)}, Y_{\theta(I_1, I_2)}, L_1 \times L_2, (\ell_{0,1}, \ell_{0,2}), C, T)$, where $C(\ell_1, \ell_2)$ is defined as $f(s)$ is defined in (11), replacing $f_i(\ell_i)$ by $C_i(\ell_i)$, and T is defined as follows:

$$T := \{((\ell_1, \ell_2), g_1 \wedge g_2, (\ell'_1, \ell'_2)) \mid (\ell_i, g_i, \ell'_i) \in T_i, \text{ for } i = 1, 2\}. \quad (12)$$

That is, M is essentially a synchronous product of M_1, M_2 .

Checking compatibility of two finite-state interfaces can be effectively done by first computing an automaton representing the composite interface $\theta(I_1, I_2)$ and then checking well-formability of the latter, using the algorithms described earlier.

A connection θ is allowed to be empty. In that case, $\rho_\theta \equiv \text{true}$, and the composition can be viewed as the *parallel composition* of two interfaces. If θ is empty, we write $I_1 \parallel I_2$ instead of $\theta(I_1, I_2)$. As may be expected, the contract of the parallel composition at a given global state is the conjunction of the original contracts at the corresponding local states, which implies that parallel composition is commutative.

LEMMA 6.4. *Consider two disjoint interfaces, $I_i = (X_i, Y_i, f_i)$, $i = 1, 2$. Then $I_1 \parallel I_2 = (X_1 \cup X_2, Y_1 \cup Y_2, f)$, where f is such that for all $s \in \mathcal{A}(X_1 \cup X_2 \cup Y_1 \cup Y_2)^*$, $f(s) \equiv f_1(s_1) \wedge f_2(s_2)$, where, for $i = 1, 2$, s_i is the projection of s to $X_i \cup Y_i$.*

A corollary of Lemma 6.4 is Theorem 6.5:

THEOREM 6.5 (COMMUTATIVITY OF PARALLEL COMPOSITION). *Let I_1 and I_2 be two disjoint interfaces. Then:*

$$I_1 \parallel I_2 = I_2 \parallel I_1.$$

THEOREM 6.6 (ASSOCIATIVITY OF CONNECTION). *Let I_1, I_2, I_3 be pairwise disjoint interfaces. Let θ_{12} be a connection between I_1, I_2 , θ_{13} a connection between I_1, I_3 , and θ_{23} a connection between I_2, I_3 . Then:*

$$(\theta_{12} \cup \theta_{13})(I_1, \theta_{23}(I_2, I_3)) = (\theta_{13} \cup \theta_{23})(\theta_{12}(I_1, I_2), I_3).$$




Fig. 4. The interface diagram of Example 6.7.

Example 6.7. Consider the diagram of stateless interfaces shown in Figure 4, where:

$$\begin{aligned} I_{id} &:= (\{x_1\}, \{y_1\}, y_1 = x_1) \\ I_{+1,2} &:= (\{x_2\}, \{y_2\}, x_2 + 1 \leq y_2 \leq x_2 + 2) \\ I_{\leq} &:= (\{z_1, z_2\}, \{\}, z_1 \leq z_2). \end{aligned}$$

This diagram can be modeled as any of the two following equivalent compositions:

$$\theta_2(I_{+1,2}, \theta_1(I_{id}, I_{\leq})) = (\theta_1 \cup \theta_2)((I_{id} \parallel I_{+1,2}), I_{\leq}),$$

where $\theta_1 := \{(y_1, z_1)\}$ and $\theta_2 := \{(y_2, z_2)\}$.

We proceed to compute the contract of the interface defined by the diagram. It is easier to consider the composition $(\theta_1 \cup \theta_2)((I_{id} \parallel I_{+1,2}), I_{\leq})$. Define $\theta_3 := \theta_1 \cup \theta_2$. From Lemma 6.4 we get:

$$I_{id} \parallel I_{+1,2} = (\{x_1, x_2\}, \{y_1, y_2\}, y_1 = x_1 \wedge x_2 + 1 \leq y_2 \leq x_2 + 2).$$

Then, for $\theta_3((I_{id} \parallel I_{+1,2}), I_{\leq})$, Formula (11) gives:

$$\Phi := (y_1 = x_1 \wedge x_2 + 1 \leq y_2 \leq x_2 + 2 \wedge y_1 = z_1 \wedge y_2 = z_2) \rightarrow z_1 \leq z_2.$$

By quantifier elimination, we get

$$\forall y_1, y_2, z_1, z_2 : \Phi \equiv x_1 \leq x_2 + 1.$$

Therefore,

$$\begin{aligned} \theta_3((I_{id} \parallel I_{+1,2}), I_{\leq}) &= (\{x_1, x_2\}, \{y_1, y_2, z_1, z_2\}, \\ & y_1 = x_1 \wedge x_2 + 1 \leq y_2 \leq x_2 + 2 \wedge z_1 \leq z_2 \\ & \wedge y_1 = z_1 \wedge y_2 = z_2 \wedge x_1 \leq x_2 + 1). \end{aligned}$$

Notice that $\text{in}(\theta_3((I_{id} \parallel I_{+1,2}), I_{\leq})) \equiv x_1 \leq x_2 + 1$. That is, because of the connection θ , new assumptions have been generated for the external inputs x_1, x_2 . These assumptions are stronger than those generated by simple composition of relations, which are $x_1 \leq x_2 + 2$ in this case.

A composite interface is not guaranteed to be well-formed, neither well-formable, even if all its components are well-formed.

Example 6.8. Consider the composite interface $\theta_3((I_{id} \parallel I_{+1,2}), I_{\leq})$ from Example 6.7, and suppose we connect its open inputs x_1, x_2 to outputs v_1, v_2 , respectively, of some other interface that guarantees $v_1 > v_2 + 1$. Clearly, the result is false, since the constraint $x_1 > x_2 + 1 \wedge x_1 \leq x_2 + 1$ is unsatisfiable.

6.2 Composition by Feedback

Our second type of composition is *feedback composition*, where an output variable of an interface I is connected to one of its input variables x . For feedback, I is required to be *Moore with respect to x* . The term “Moore interfaces” has been introduced

in Chakrabarti et al. [2002]. Our definition is similar in spirit, but less restrictive than the one in Chakrabarti et al. [2002]. Both definitions are inspired by Moore machines, where the outputs are determined by the current state alone and do not depend directly on the input.

In our version, an interface is Moore with respect to a given input variable x , meaning that the contract may depend on the current state as well as on input variables other than x . This allows to connect an output to x to form a feedback loop without creating causality cycles.

Definition 6.9 (Moore interfaces). An interface $I = (X, Y, f)$ is called *Moore with respect to* $x \in X$ iff for all $s \in f$, $f(s)$ is a property over $(X \cup Y) \setminus \{x\}$. I is called simply *Moore* when it is Moore with respect to every $x \in X$.

Note that a source interface is by definition Moore, since it has no input variables. Note also that although the contract of a Moore interface should not depend on the *current* value of an input variable, it may very well depend on *past* values of such a variable, which influence the state s . An example where this occurs is the unit delay.

Example 6.10 (Unit delay). A *unit delay* is a basic building block in many modeling languages (including Simulink and SCADE). Its specification is roughly: “output at time k the value of the input at time $k - 1$; at time $k = 0$ (initial time), output some initial value v_0 .” We can capture this specification as an interface:

$$I_{ud} := (\{x\}, \{y\}, f_{ud}),$$

where f_{ud} is defined as follows:

$$\begin{aligned} f_{ud}(\varepsilon) &:= (y = v_0) \\ f_{ud}(s \cdot a) &:= (y = a(x)). \end{aligned}$$

That is, initially the contract guarantees $y = v_0$. Then, when the state is some sequence $s \cdot a$, the contract guarantees $y = a(x)$, where $a(x)$ is the last value assigned to input x . I_{ud} is Moore (with respect to its unique input variable) since all its contracts are properties over y only.

Definition 6.11 (Composition by feedback). Let $I = (X, Y, f)$ be an interface. A *feedback connection* κ on I is a pair $(y, x) \in Y \times X$. κ is *valid* if I is Moore with respect to x . Define $\rho_\kappa := (x = y)$. A valid feedback connection κ defines the interface:

$$\kappa(I) := (X \setminus \{x\}, Y \cup \{x\}, f_\kappa) \tag{13}$$

$$f_\kappa(s) := f(s) \wedge \rho_\kappa, \quad \text{for all } s \in \mathcal{A}(X \cup Y)^*. \tag{14}$$

In the sequel, when we talk about feedback connections we implicitly assume they are valid.

For finite-state interfaces, feedback is computable. Let $M = (X, Y, L, \ell_0, C, T)$ be a finite-state automaton representing I . First, to check whether M represents a Moore interface w.r.t. a given input variable $x \in X$, it suffices to make sure that for every location $\ell \in L$, $C(\ell)$ does not refer to x . Then, if $\kappa = (y, x)$, the interface $\kappa(I)$ can be




Fig. 5. An interface diagram with feedback.

represented as $M' := (X \setminus \{x\}, Y \cup \{x\}, L, \ell_0, C', T)$, where $C'(\ell) := C(\ell) \wedge x = y$, for all $\ell \in L$.

THEOREM 6.12 (COMMUTATIVITY OF FEEDBACK). *Let $I = (X, Y, f)$ be an interface which is Moore with respect to both $x_1, x_2 \in X$, where $x_1 \neq x_2$. Let $\kappa_1 = (y_1, x_1)$ and $\kappa_2 = (y_2, x_2)$ be feedback connections. Then κ_1 is valid for both I and $\kappa_2(I)$, κ_2 is valid for both I and $\kappa_1(I)$, and*

$$\kappa_1(\kappa_2(I)) = \kappa_2(\kappa_1(I)).$$

Let K be a set of feedback connections, $K = \{\kappa_1, \dots, \kappa_n\}$, such that $\kappa_i = (y_i, x_i)$, and all x_i are pairwise distinct, for $i = 1, \dots, n$. Let I be an interface that is Moore with respect to all x_1, \dots, x_n . We denote by $K(I)$ the interface $\kappa_1(\kappa_2(\dots \kappa_n(I) \dots))$. By commutativity of feedback composition, the resulting interface is independent from the order of application of feedback connections. We will use the notation $\text{InVars}(K) := \{x_i \mid (y_i, x_i) \in K\}$, for the set of input variables connected in K .

THEOREM 6.13 (COMMUTATIVITY BETWEEN CONNECTION AND FEEDBACK). *Let I_1, I_2 be disjoint interfaces and let θ be a connection between I_1, I_2 . Let κ_1, κ_2 be valid feedback connections on I_1, I_2 , respectively. Suppose that $\text{InVars}(\kappa_2) \cap \text{InVars}(\theta) = \emptyset$. Then:*

$$\kappa_1(\theta(I_1, I_2)) = \theta(\kappa_1(I_1), I_2) \quad \text{and} \quad \kappa_2(\theta(I_1, I_2)) = \theta(I_1, \kappa_2(I_2)).$$

THEOREM 6.14 (PRESERVATION OF MOORENESS BY CONNECTION). *Let I_1, I_2 be disjoint interfaces such that $I_i = (X_i, Y_i, f_i)$, for $i = 1, 2$. Let θ be a connection between I_1, I_2 .*

- (1) *If I_1 is Moore w.r.t. $x_1 \in X_1$, then $\theta(I_1, I_2)$ is Moore w.r.t. x_1 .*
- (2) *If I_1 is Moore and $\text{InVars}(\theta) = X_2$, then $\theta(I_1, I_2)$ is Moore.*
- (3) *If I_2 is Moore w.r.t. $x_2 \in X_2$ and $x_2 \notin \text{InVars}(\theta)$, then $\theta(I_1, I_2)$ is Moore w.r.t. x_2 .*

An interesting question is to what extent and how to transform a given diagram of interfaces, such as the one shown in Figure 5, to a valid expression of interface compositions. This cannot be done for arbitrary diagrams, due to restrictions on feedback, but it can be done for diagrams that satisfy the following condition: every dependency cycle in the diagram, formed by block connections, must visit at least one input variable x of some interface I , such that I is Moore w.r.t. x . If this condition holds, then we say that the diagram is *causal*. For example, the diagram in Figure 5 is causal iff I_1 is Moore w.r.t. x_2 or I_2 is Moore w.r.t. x_4 .

We can systematically transform causal interface diagrams into expressions of interface compositions as follows. First, we remove from the diagram any *Moore connections*. A connection from output variable y to input variable x is a Moore connection if the interface I where x belongs to is Moore w.r.t. x . Because the original diagram is by hypothesis causal, the diagram obtained after removing Moore connections is guaranteed to have no dependency cycles. This acyclic diagram can be easily transformed into

an expression involving only interface compositions by connection. By associativity of connection (Theorem 6.6), the order in which these connections are applied does not matter. Call the resulting interface I_c . Then, the removed Moore connections can be turned into feedback compositions, and applied to I_c . Because Mooreness is preserved by connection (Theorem 6.14), I_c is guaranteed to be Moore w.r.t. any input variable x that is the destination of a Moore connection. Therefore, the above feedback compositions are valid for I_c . Moreover, because of commutativity of feedback (Theorem 6.12), the resulting interface is again uniquely defined.

Example 6.15. Consider the diagram of interfaces shown in Figure 5. Suppose that I_1 is Moore with respect to x_2 . Then, the diagram can be expressed as any of the two compositions

$$\kappa\left(\theta_1(I_1, (I_2 \parallel I_3))\right) = \theta_3\left(\kappa(\theta_2(I_1, I_2)), I_3\right),$$

where $\theta_1 := \{(y_1, x_4), (y_2, x_3)\}$, $\theta_2 := \{(y_1, x_4)\}$, $\theta_3 := \{(y_2, x_3)\}$, and $\kappa := (y_4, x_2)$. The two expressions are equivalent, since, by Theorem 6.13, $\theta_3\left(\kappa(\theta_2(I_1, I_2)), I_3\right) = \kappa\left(\theta_3(\theta_2(I_1, I_2), I_3)\right)$, and by Theorem 6.6, $\theta_3(\theta_2(I_1, I_2), I_3) = \theta_1(I_1, (I_2 \parallel I_3))$.

LEMMA 6.16. *Let $I = (X, Y, f)$ be a Moore interface with respect to $x \in X$. Let $\kappa = (y, x)$ be a feedback connection on I . Let $\kappa(I) = (X \setminus \{x\}, Y \cup \{y\}, f_\kappa)$. Then:*

- (1) $f_\kappa \subseteq f$.
- (2) For any $s \in f_\kappa$, $\text{in}(f_\kappa(s)) \equiv \text{in}(f(s))$.

THEOREM 6.17 (FEEDBACK PRESERVES WELL-FORMEDNESS). *Let I be a Moore interface with respect to some of its input variables, and let κ be a valid feedback connection on I . If I is well-formed, then $\kappa(I)$ is well-formed.*

Feedback does not preserve well-formability.

Example 6.18. Consider a finite-state interface I_f with two states, s_0 (the initial state) and s_1 , one input variable x and one output variable y . I_f remains at state s_0 when $x \neq 0$ and moves from s_0 to s_1 when $x = 0$. Let $\phi_0 := y = 0$ be the contract at state s_0 and let $\phi_1 := \text{false}$ be the contract at state s_1 . I_f is not well-formed because ϕ_1 is unsatisfiable while state s_1 is reachable. I_f is well-formable, however: it suffices to restrict ϕ_0 to $\phi'_0 := y = 0 \wedge x \neq 0$. Denote the resulting (well-formed) interface by I'_f . Note that I_f is Moore with respect to x , whereas I'_f is not. Let κ be the feedback connection (y, x) . Because I_f is Moore, $\kappa(I_f)$ is defined, and is such that its contract at state s_0 is $y = 0 \wedge x = y$, and its contract at state s_1 is $\text{false} \wedge x = y \equiv \text{false}$. $\kappa(I_f)$ is not well-formable: indeed, $y = 0 \wedge x = y$ implies $x = 0$, therefore, state s_1 cannot be avoided.

7. HIDING

As can be seen in Example 6.7, composition often creates redundant output variables, in the sense that some of these variables are equal to each other. This happens because input variables that get connected become output variables. To remove redundant output variables, we propose a *hiding* operator. Hiding may also be used to remove other output variables that may not be redundant, provided they do not influence the evolution of contracts, as we shall see below.

For a stateless interface $I = (X, Y, \phi)$, the (stateless) interface resulting from hiding an output variable $y \in Y$ can simply be defined as:

$$\text{hide}(y, I) := (X, Y \setminus \{y\}, \exists y : \phi).$$

This definition does not directly extend to the general case of stateful interfaces, however. The reason is that the contract of a stateful interface I may depend on the history of y . Then, hiding y is problematic because it results in the environment not being able to uniquely determine the contract based on the history of observations. This results in particular in refinement not being preserved by hiding, as we show later in Example 9.13. To avoid these problems, we allow hiding only for those outputs which do not influence the evolution of the contract.

Given $s, s' \in \mathcal{A}(X \cup Y)^*$ such that $|s| = |s'|$ (i.e., s, s' have same length), and given $Z \subseteq X \cup Y$, we say that s and s' *agree on Z* , denoted $s =_Z s'$, when for all $i \in \{1, \dots, |s|\}$, and all $z \in Z$, $s_i(z) = s'_i(z)$. Given interface $I = (X, Y, f)$, we say that f is *independent from z* if for every $s, s' \in f$, $s =_{(X \cup Y) \setminus \{z\}} s'$ implies $f(s) = f(s')$. That is, the evolution of z does not affect the evolution of f .

Notice that f being independent from z does *not* imply that f cannot refer to variable z . Indeed, all stateless interfaces trivially satisfy the independence condition: their contracts are invariant in time, that is, they do not depend on the evolution of variables. Clearly, the contract of a stateless interface can refer to any of its variables. Conversely, even if the contracts specified by f do not refer to z , f may still depend on z , because the evolution of contracts may depend on z . For example, suppose that $f(\varepsilon) \equiv \text{true}$, and that $f(z = 0)$ is different from $f(z = 1)$, although no contract refers to z . Here, $f(z = k)$ denotes the contract at a state where $z = k$. In this case, f depends on z since the value z assumes at the first round determines the contract to be used in the second round.

The above notion of independence is weaker than redundancy in variables, as we show next. First, we formalize redundancy in variables. Given $z \in X \cup Y$, we say that z is *redundant in f* if there exists $z' \in X \cup Y$ such that $z' \neq z$, and for all $s \in f$, for all $i \in \{1, \dots, |s|\}$, $s_i(z) = s_i(z')$. It should be clear that all outputs in $\text{InVars}(\theta)$ in an interface obtained by connection θ are redundant (see Definition 6.1). Similarly, in an interface obtained by feedback $\kappa = (y, x)$, newly introduced output variable x is redundant (see Definition 6.11).

LEMMA 7.1. *If z is redundant in f , then f is independent from z .*

When f is independent from z , f can be viewed as a function from $\mathcal{A}((X \cup Y) \setminus \{z\})^*$ to $\mathcal{F}(X \cup Y)$ instead of a function from $\mathcal{A}(X \cup Y)^*$ to $\mathcal{F}(X \cup Y)$. We use this when we write $f(s)$ for $s \in \mathcal{A}((X \cup Y) \setminus \{z\})^*$ in the following definition.

Definition 7.2 (Hiding). Let $I = (X, Y, f)$ be an interface and let $y \in Y$, such that f is independent from y . Then $\text{hide}(y, I)$ is defined to be the interface

$$\text{hide}(y, I) := (X, Y \setminus \{y\}, f'), \quad (15)$$

such that for any $s \in \mathcal{A}(X \cup Y \setminus \{y\})^*$, $f'(s) := \exists y : f(s)$.

For finite-state interfaces, hiding is computable. Let $M = (X, Y, L, \ell_0, C, T)$ be a finite-state automaton representing I . We first need to ensure that the contract of I is independent from y . A simple way to do this is to check that no guard of M refers to y . This condition is sufficient, but not necessary. Consider, for example, two complementary guards $y < 1$ and $y \geq 1$ whose transitions lead to locations with identical contracts. Then the two locations may be merged to a single one, and the two transitions to a single transition with guard true. Another situation where the above condition may be too strict is when a guard refers to y but y is redundant. In that case, all occurrences of y in guards of M can be replaced by its equal variable y' . Once independence from y is ensured, $\text{hide}(y, I)$ can be represented as $M' := (X, Y \setminus \{y\}, L, \ell_0, C', T)$, where $C'(\ell) := \exists y : C(\ell)$, for all $\ell \in L$.




Fig. 6. Illustration of pluggability.

8. ENVIRONMENTS, PLUGGABILITY AND SUBSTITUTABILITY

We wish to formalize the notion of interface contexts and substitutability, and we introduce *environments* for that purpose. Environments are interfaces. An interface I can be connected to an environment E to form a closed-loop system, as illustrated in Figure 6. E acts both as a *controller* and an *observer* for I . It is a controller in the sense that it “steers” I by providing inputs to it, depending on the outputs it receives. At the same time, E acts as an observer, that monitors the inputs consumed and outputs produced by I , and checks whether a given property is satisfied. These notions are formalized in Definition 8.1 that follows.

Before giving the definition, however, a remark is in order. Interfaces and environments are to be connected in a closed loop, as illustrated in Figure 6. In order to do this in our setting, every dependency cycle must be “broken” by a Moore connection, as prescribed by the transformation of interface diagrams to composition expressions, given in Section 6.2. It can be seen that, in the case of two interfaces connected in closed-loop, the above requirement implies that one of the two interfaces is Moore. For instance, consider Figure 6. If I is not Moore w.r.t. x_2 , then E must be Moore w.r.t. to both \hat{y}_1 and \hat{y}_2 , so that both feedback connections can be formed. Similarly, if E is not Moore w.r.t. \hat{y}_2 , say, then I must be Moore w.r.t. both x_1, x_2 . This remark justifies the following definition.

Definition 8.1 (Environments and pluggability). Consider interfaces $I = (X, Y, f)$ and $E = (\hat{Y}, \hat{X}, f_e)$. E is said to be an *environment* for I if there exist bijections between X and \hat{X} , and between Y and \hat{Y} . \hat{X} are called the *mirror variables* of X , and similarly for \hat{Y} and Y . For $x \in X$, we denote by \hat{x} the corresponding (by the bijection) variable in \hat{X} , and similarly with y and \hat{y} . I is said to be *pluggable* to E , denoted $I \rightleftharpoons E$, iff the following conditions hold.

- I is Moore or E is Moore.
- If E is Moore, then the interface $K(\theta(E, I))$ is well-formed, where $\theta := \{(\hat{x}, x) \mid x \in X\}$ and $K := \{(y, \hat{y}) \mid y \in Y\}$. Notice that, because E is Moore and $\text{InVars}(\theta) = X$, part 2 of Theorem 6.14 applies, and guarantees that $\theta(E, I)$ is Moore. Therefore, $K(\theta(E, I))$ is well-defined.
- If I is Moore, then the interface $K(\theta(I, E))$ is well-formed, where $\theta := \{(y, \hat{y}) \mid y \in Y\}$ and $K := \{(\hat{x}, x) \mid x \in X\}$.

Note that, by definition, I is pluggable to E iff E is pluggable to I .

Example 8.2. Consider interfaces I_1 and I_2 from Example 5.6 and environments E_1, E_2, E_3 of Figure 7 (implicitly, transitions without guards are assumed to have guard true). It can be checked that both I_1 and I_2 are pluggable to E_1 . I_1 is not pluggable to neither E_2 nor E_3 : indeed, the output guarantee $\hat{x} \geq 0$ of these two environments is not strong enough to meet the input assumption $x > 0$ of I_1 . I_2 is not pluggable to E_2 : although the input assumption of I_2 is true, I_2 guarantees $y > 0$ only when $x > 0$. Therefore, the guard $\hat{y} \leq 0$ of E_2 is enabled in some cases,





Fig. 7. Three environments.

Fig. 8. A Moore interface I and a non-Moore environment E .

leading to location with contract `false`, which means that the closed-loop interface is not well-formed. On the other hand, I_2 is pluggable to E_3 .

THEOREM 8.3 (PLUGGABILITY AND WELL-FORMABILITY).

- If an interface I is well-formable, then there exists an environment E for I such that $I \sqsubseteq E$.
- If there exists an environment E for interface I such that $I \sqsubseteq E$ and I is not Moore, then I is well-formable.

Example 8.4. Consider interfaces I and E shown in Figure 8. Observe that I is Moore and $I \sqsubseteq E$. However, I is not well-formable.

Example 8.4 shows that the non-Mooreness assumption on I is indeed necessary in part 2 of Theorem 8.3. This example also illustrates an aspect of our definition of well-formability, which may appear inappropriate for Moore interfaces: indeed, interface I of Figure 8 is non-well-formable, yet there is clearly an environment that can be plugged to I so that `false` location is avoided. An alternative definition of well-formability for an interface I would have been existence of an environment that can be plugged to I . This would make Theorem 8.3 a tautology. Nevertheless, we opt for Definition 5.11, which allows to transform interfaces into a “canonical form” where all contracts are satisfiable.

Definition 8.5 (Substitutability). We say that interface I' may replace interface I (or I' may be substituted for I), denoted $I \rightarrow_e I'$, iff for any environment E , if I is pluggable to E then I' is pluggable to E . We say that I and I' are *mutually substitutable*, denoted $I \equiv_e I'$, iff both $I \rightarrow_e I'$ and $I' \rightarrow_e I$ hold.

As we shall show in Theorem 9.16, for well-formed interfaces, mutual substitutability coincides with interface equality.

9. REFINEMENT

Definition 9.1 (Refinement). Consider two interfaces $I = (X, Y, f)$ and $I' = (X', Y', f')$. We say that I' *refines* I , written $I' \sqsubseteq I$, iff $X' = X$, $Y' = Y$, and for any $s \in f \cap f'$, the following formula is valid:

$$\text{in}(f(s)) \rightarrow \left(\text{in}(f'(s)) \wedge (f'(s) \rightarrow f(s)) \right) \quad (16)$$

Condition 16 can be rewritten equivalently as the conjunction of the following two conditions:

$$\begin{aligned} \text{in}(f(s)) &\rightarrow \text{in}(f'(s)) && (17) \\ (\text{in}(f(s)) \wedge f'(s)) &\rightarrow f(s) && (18) \end{aligned}$$

Condition (17) states that every input assignment that is legal in I is also legal in I' . This guarantees that, for any possible input assignment that can be provided to I by a context C , if this assignment is accepted by I then it is also accepted by I' . Condition (18) states that, for every input assignment that is legal in I , all output assignments that can be possibly produced by I' from that input, can also be produced by I . This guarantees that if C accepts the assignments produced by I then it also accepts those produced by I' .

It should be noted that the refinement conditions are required only for states that belong in both f and f' . The intuition for this choice is as follows. The initial state is ε and by definition $\varepsilon \in f \cap f'$. At this state, we only wish to consider legal inputs for I , that is, $a_X \in \text{in}(f(\varepsilon))$. Otherwise, I' is free to behave as it wishes since the behavior is not possible in I . Condition (17) then implies that $a_X \in \text{in}(f'(\varepsilon))$. Next, we wish to consider only outputs that I' may produce given a_X , that is, a_Y such that $(a_X, a_Y) \in f'(\varepsilon)$. Otherwise, I is free to behave as it wishes, since the behavior is not possible in I' . Condition (18) then implies that $(a_X, a_Y) \in f(\varepsilon)$. Therefore, $(a_X, a_Y) \in f \cap f'$, that is, the requirements should be applied only to states of length one that belong in both f and f' . Reasoning inductively, the same can be derived for states of arbitrary length.

A remark is in order regarding the constraint $X' = X$ and $Y' = Y$ imposed during refinement. This constraint may appear as too strict, but we argue that it is not. To begin, recall that $I' \sqsubseteq I$ should imply that I' can replace I in any context. In our setting, contexts are formalized as environments. Consider such an environment with controller C . C provides values to the input variables of I , and requires values from the output variables of I . Suppose I' has an input variable x that I does not have, that is, there exists $x \in X' \setminus X$. In general, C may not provide x . In that case, I' cannot replace I , because by doing so, input x would remain free. Therefore, $X' \subseteq X$ must hold. Similarly, suppose that there exists $y \in Y \setminus Y'$. In general, C may require y , that is, y may be a free input for C . In that case, I' cannot replace I , because by doing so, y would remain free. Therefore, $Y \subseteq Y'$ must hold.

Now, suppose that X' is a strict subset of X or Y' is a strict superset of Y (or both). Then, we can easily modify I and I' as follows: we add to X' all the input variables missing from I , so that $X' = X$, and we add to Y all the output variables missing from I' , so that $Y = Y'$. While doing so, we do not change the contracts of either I or I' : the contracts simply ignore the additional variables, that is, do not impose any constraints on their values. It can be seen that this transformation preserves the validity of refinement Condition 16. Indeed, $\text{in}(\phi) \rightarrow (\text{in}(\phi') \wedge (\phi' \rightarrow \phi))$ holds when ϕ is over $X \cup Y$ and ϕ' is over $X' \cup Y'$ iff it holds when both ϕ and ϕ' are taken to be over $X \cup Y'$, provided $X' \subseteq X$ and $Y' \supseteq Y$. Therefore, without loss of generality, we require $X = X'$ and $Y = Y'$.

Example 9.2 (Buffer interface refinements). This example builds on Example 5.8. Consider Figure 9. It depicts a variant of the single-place buffer interface, where the buffer may fail to complete a read or write operation. This interface has one more boolean output variable, namely, *ack*, in addition to those of Example 5.8, and two more locations, *after_read* and *after_write*. Its global contract is identical to that of Example 5.8. So are local contracts at locations q_0 and q_1 . After a write operation, the

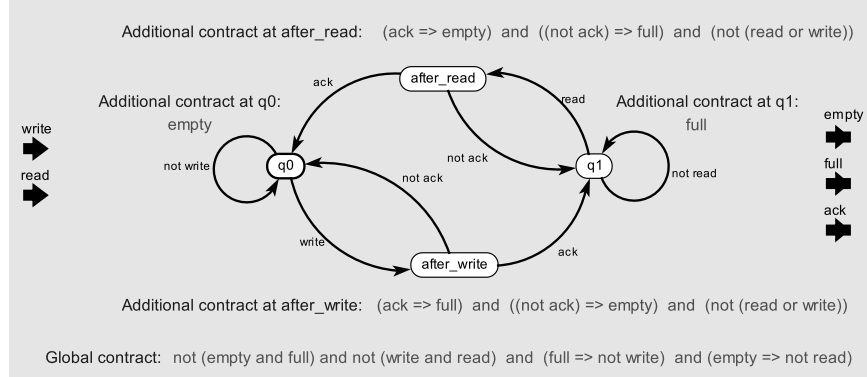
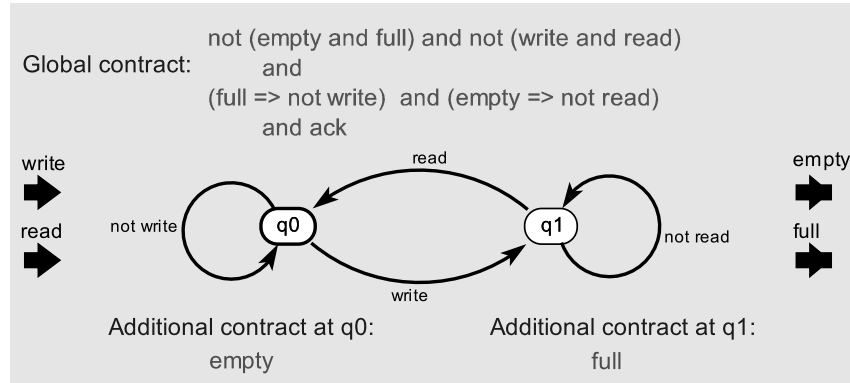


Fig. 9. Interface for a buffer of size 1 that may fail to do a read or write.

Fig. 10. Buffer interface of Figure 2 with additional output variable *ack*.

interface moves to location *after_write*, where it non-deterministically chooses to set *ack* to true or false: setting it to true means the write was successful, false means the write failed. The meaning is symmetric for read. This particular interface does not allow read or write operations in the two intermediate locations.

It is natural to expect that a buffer that never fails can replace a buffer that may fail. We would like to have a formal guarantee of this, in terms of refinement of their corresponding interfaces. That is, we would like the interface of Figure 2 to refine the one of Figure 9. This does not immediately hold, since *ack* is not a variable of the former. We can easily add it however, obtaining the interface shown in Figure 10. This buffer never fails, therefore, *ack* is always true. With this modification, the interface of Figure 10 refines the one of Figure 9. On the other hand, the converse is not true: the interface of Figure 9 does not refine the one of Figure 10, because in the latter output *ack* is always true, whereas in the former it can also be false.

With respect to the above discussion on the $X = X'$ and $Y = Y'$ requirements, note that in this example the condition $X' \subseteq X$ and $Y \subseteq Y'$ does not hold: indeed, Y (the outputs of Figure 9) includes *ack* whereas Y' (the outputs of Figure 2) does not. For this reason, *ack* is not simply a “dummy” variable in this case, and we need to specify a contract for it, as done in the revised interface of Figure 10.

This example also illustrates the fact that our notion of refinement is different from language inclusion. For instance, the sequence

$$(empty, \neg full, ack, write, \neg read) \cdot (\neg empty, full, ack, \neg write, read)$$

belongs in the language (i.e., contract) of the interface of Figure 10, but not in the language of the interface of Figure 9. This is because the latter does not allow a “read” to happen at state “after_write”.

For finite-state interfaces, refinement can be checked as follows. Let $M_i = (X, Y, L_i, \ell_{0,i}, C_i, T_i)$ be finite-state automata representing I_i , for $i = 1, 2$, respectively. We first build a synchronous product

$$M := (X, Y, L_1 \times L_2 \cup \{\ell_{good}, \ell_{bad}\}, (\ell_{0,1}, \ell_{0,2}), C, T),$$

where $C(\ell_1, \ell_2) := \text{in}(C_1(\ell_1))$ for all $(\ell_1, \ell_2) \in L_1 \times L_2$, $C(\ell_{good}) := \text{true}$, $C(\ell_{bad}) := \text{false}$, and:

$$T := \{((\ell_1, \ell_2), g_{both} \wedge g_1 \wedge g_2, (\ell'_1, \ell'_2)) \mid (\ell_i, g_i, \ell'_i) \in T_i, \text{ for } i = 1, 2\} \\ \cup \{((\ell_1, \ell_2), g_{bad}, \ell_{bad}), ((\ell_1, \ell_2), g_{good}, \ell_{good}), (\ell_{good}, \text{true}, \ell_{good})\} \quad (19)$$

$$g_{both} := C_1(\ell_1) \wedge C_2(\ell_2) \quad (20)$$

$$g_{good} := \text{in}(C_1(\ell_1)) \wedge \text{in}(C_2(\ell_2)) \wedge \neg C_2(\ell_2) \quad (21)$$

$$g_{bad} := \text{in}(C_1(\ell_1)) \wedge (\neg \text{in}(C_2(\ell_2)) \vee (C_2(\ell_2) \wedge \neg C_1(\ell_1))). \quad (22)$$

Notice that guard g_{bad} encodes the negation of the refinement Condition (16). Also note that $g_{both}, g_{good}, g_{bad}$ are pairwise disjoint, and such that $g_{both} \vee g_{good} \vee g_{bad} \equiv \text{in}(C_1(\ell_1))$, for all $(\ell_1, \ell_2) \in L_1 \times L_2$. This ensures determinism of M . It can be checked that $I_2 \sqsubseteq I_1$ iff location ℓ_{bad} is unreachable.

9.1 Properties of the Refinement Relation

We proceed to state the main properties of refinement. First, observe that, perhaps surprisingly, interfaces with false contracts (i.e., $f = \{\varepsilon\}$) are “top” elements with respect to the \sqsubseteq order, that is, they are refined by any interface that has the same input and output variables. This is in accordance with the spirit of refinement as a condition for substitutability. The false interface is not pluggable to any environment; therefore, it can be replaced by any interface.

We next provide a result used in the proof of the theorems that follow.

LEMMA 9.3. *Let $I = (X, Y, f)$, $I' = (X, Y, f')$, $I'' = (X, Y, f'')$ be interfaces such that $I'' \sqsubseteq I'$ and $I' \sqsubseteq I$. Then $f \cap f'' \sqsubseteq f'$.*

An illustration of the preceding lemma can be found in the division example in Section 2.3, where $\phi_1 \wedge \phi_4 \equiv \phi_2$.

THEOREM 9.4 (PARTIAL ORDER). \sqsubseteq is a partial order, that is, a reflexive, antisymmetric, and transitive relation.

THEOREM 9.5. *Let I, I' be stateless interfaces such that $I' \sqsubseteq I$. If I is well-formed, then I' is well-formed.*

Theorem 9.5 does not generally hold for stateful interfaces: the reason is that, because I' may accept more inputs than I , there may be states that are reachable in I' but not in I , and the contract of I' in these states may be unsatisfiable. When this situation does not occur, refinement preserves well-formedness also in the stateful case. Moreover, refinement always preserves well-formability.

THEOREM 9.6 (REFINEMENT AND WELL-FORMEDNESS/-FORMABILITY). *Let I, I' be interfaces such that $I' \sqsubseteq I$.*

- (1) *If I is well-formed and $f(I') \subseteq f(I)$, then I' is well-formed.*
- (2) *If I, I' are sources and I is well-formed, then I' is also well-formed.*
- (3) *If I is well-formable, then I' is well-formable.*

The following lemma is used in the proof of Theorem 9.8.

LEMMA 9.7. *Consider two disjoint interfaces I_1 and I_2 , and a connection θ between I_1, I_2 . Let f_1 and f_2 be the projections of $f(\theta(I_1, I_2))$ to states over the variables of I_1 and I_2 , respectively. Then $f_1 \subseteq f(I_1)$ and $f_2 \subseteq f(I_2)$.*

Theorems 9.8 and 9.9 that follow state a major property of our theory, namely, that refinement is preserved by composition.

THEOREM 9.8 (CONNECTION PRESERVES REFINEMENT). *Consider two disjoint interfaces I_1 and I_2 , and a connection θ between I_1, I_2 . Let I'_1, I'_2 be interfaces such that $I'_1 \sqsubseteq I_1$ and $I'_2 \sqsubseteq I_2$. Then $\theta(I'_1, I'_2) \sqsubseteq \theta(I_1, I_2)$.*

Notice that Theorem 9.8 holds independently of whether the connection yields a well-formed interface or not, that is, independently of whether the composed interfaces are compatible. This is a reason why we do not impose compatibility as a condition for composition, as we mentioned earlier. Together with Theorems 9.5 and 9.6, Theorem 9.8 guarantees that if the refined composite interface is well-formed/formable, then so is the refining one. In particular, if I_1 and I_2 are compatible with respect to θ , then so are I'_1 and I'_2 .

THEOREM 9.9 (FEEDBACK PRESERVES REFINEMENT). *Let I, I' be interfaces such that $I' \sqsubseteq I$. Suppose both I and I' are Moore interfaces with respect to one of their input variables, x . Let $\kappa = (y, x)$ be a feedback connection. Then $\kappa(I') \sqsubseteq \kappa(I)$.*

Note that the assumption that I' be Moore w.r.t. x in Theorem 9.9 is essential. Indeed, Mooreness is not generally preserved by refinement.

Example 9.10. Consider the stateless interfaces $I_{\text{even}} := (\{x\}, \{y\}, y \div 2 = 0)$, where \div denotes the modulo operator, and $I_{\times 2} := (\{x\}, \{y\}, y = 2x)$. I_{even} is Moore. $I_{\times 2}$ is not Moore. Yet $I_{\times 2} \sqsubseteq I_{\text{even}}$.

It is instructive at this point to justify our restrictions regarding feedback composition, by illustrating some of the problems that would arise if we allowed arbitrary feedback.

Example 9.11. This example is borrowed from Doyen et al. [2008]. Suppose I_{true} is an interface on input x and output y , with trivial contract true, making no assumptions on the inputs and no guarantees on the outputs. Suppose $I_{y \neq x}$ is another interface on x and y , with contract $y \neq x$, meaning that it guarantees that the value of the output will be different from the value of the input. As expected, $I_{y \neq x}$ refines I_{true} : because $I_{y \neq x}$ is “more deterministic” than I_{true} , that is, the output guarantees of $I_{y \neq x}$ are stronger. Now, consider the feedback connection $x = y$. This could be considered an allowed connection for I_{true} , since it does not contradict its contract: the resulting interface would be $I_{x=y}$ with contract $x = y$. But the same feedback connection contradicts the contract of $I_{y \neq x}$: the resulting interface would be I_{false} with contract false. Although $I_{y \neq x}$ refines I_{true} , I_{false} does not refine $I_{x=y}$, therefore, allowing arbitrary feedback would violate preservation of refinement by feedback. Notice that both I_{true}




Fig. 11. Example illustrating the need for independence from hidden variables.

and $I_{y \neq x}$ are input-complete, which means that this problem is present also in that special case.

THEOREM 9.12 (HIDING PRESERVES REFINEMENT). *Let $I_1 = (X, Y, f_1)$ and $I_2 = (X, Y, f_2)$ be two interfaces such that $I_2 \sqsubseteq I_1$. Let $y \in Y$ be such that both f_1 and f_2 are independent from y . Then $\text{hide}(y, I_2) \sqsubseteq \text{hide}(y, I_1)$.*

It is worth noting that Theorem 9.12 would not hold if we were to define hiding without requiring independence of contracts from hidden variables. The example that follows illustrates this.

Example 9.13. Consider the interfaces shown in Figure 11. I_1 and I_2 have a single input variable x and a single output y . It can be verified that $I_2 \sqsubseteq I_1$. I_2 is independent from y , whereas I_1 is not. Therefore, $\text{hide}(y, I_2)$ is defined (and shown in the figure), whereas $\text{hide}(y, I_1)$ is not defined. Suppose we were to define the latter as interface I_1' shown in the figure, which corresponds to existentially quantifying away y from all contracts, as is usually done. Then hiding would not preserve refinement. Indeed, $\text{hide}(y, I_2) \not\sqsubseteq I_1'$, because $x \cdot \neg x$ is a legal input sequence in I_1' but not in $\text{hide}(y, I_2)$.

THEOREM 9.14 (REFINEMENT AND SUBSTITUTABILITY). *Let I, I' be two interfaces.*

- (1) *If $I' \sqsubseteq I$, then I' can replace I .*
- (2) *If $I' \not\sqsubseteq I$, and I is well-formed, then I' cannot replace I .*

The requirement that I be well-formed in part 2 of Theorem 9.14 is necessary, as the following example shows.

Example 9.15. Consider the finite-state interfaces I and I' defined by the automata shown in Figure 3. Both have a single boolean input variable x . I' is well-formed but I is not (I is well-formable, however, and I' is a witness). $I' \not\sqsubseteq I$, because at the initial state the input $x = \text{false}$ is legal for I but not for I' . But there is no environment E such that $I \models E$ but $I' \not\models E$.

We next state a result that is not about refinement, but follows from properties of refinement.

THEOREM 9.16. *Let I, I' be well-formed interfaces. Then $I \equiv_e I'$ iff $I = I'$.*

PROOF. By Theorem 9.14, $I \equiv_e I'$ implies $I' \sqsubseteq I$ and $I \sqsubseteq I'$. The result follows by antisymmetry of refinement (Theorem 9.4). \square

9.2 Discussion: Alternative Definition of Refinement

The reader may wonder why Condition (18) could not be replaced with the simpler condition:

$$f'(s') \rightarrow f(s). \quad (23)$$

Indeed, for input-complete interfaces, Condition (16) reduces to Condition (23); see Theorem 11.8 in Section 11. In general, however, Condition (16) is too strong in the sense that it results in a refinement condition that is sufficient but not necessary for substitutability, as the following example demonstrates.

Example 9.17. Consider interface $I_{id} := (\{x\}, \{y\}, x = y)$, and interface I_1 from Example 5.6. It can be checked that $I_{id} \sqsubseteq I_1$. If we used Condition (23) instead of Condition (18) in the definition of refinement, then I_{id} would not refine I_1 : this is because $x = y \not\rightarrow x > 0$. Yet, by Theorem 9.14, I_{id} can replace I_1 , that is, there is no environment E such that $I_1 \models E$ but $I_{id} \not\models E$.

10. SHARED REFINEMENT AND SHARED ABSTRACTION

A *shared refinement* operator \sqcap is introduced in Doyen et al. [2008] for A/G interfaces, as a mechanism to combine two such interfaces I and I' into a single interface $I \sqcap I'$ that refines both I and I' : $I \sqcap I'$ is able to accept inputs that are legal in either I or I' , and provide outputs that are legal in both I and I' . Because of this, $I \sqcap I'$ can replace both I and I' , which, as argued in Doyen et al. [2008], is important for component reuse. A similar mechanism called *fusion* has also been proposed in Benveniste et al. [2008].

Doyen et al. [2008] also discuss shared refinement for extended (i.e., relational) interfaces and conjectures that it represents the greatest lower bound with respect to refinement. We show that this holds only if a certain condition is imposed. We call this condition *shared refinability*. It states that for every inputs that is legal in both I and I' , the corresponding sets of outputs of I and I' must have a nonempty intersection. Otherwise, it is impossible to provide an output that is legal in both I and I' .

Definition 10.1 (Shared refinement). Two interfaces $I = (X, Y, f)$ and $I' = (X', Y', f')$ are *shared-refinable* if $X = X'$, $Y = Y'$ and the following formula is true for all $s \in f \cap f'$:

$$\forall X : (\text{in}(f(s)) \wedge \text{in}(f'(s))) \rightarrow \exists Y : (f(s) \wedge f'(s)). \quad (24)$$

In that case, the *shared refinement* of I and I' , denoted $I \sqcap I'$, is the interface defined as follows:

$$\begin{aligned} I \sqcap I' &:= (X, Y, f_{\sqcap}) \\ f_{\sqcap}(s) &:= (\text{in}(f(s)) \vee \text{in}(f'(s))) \wedge (\text{in}(f(s)) \rightarrow f(s)) \wedge (\text{in}(f'(s)) \rightarrow f'(s)). \end{aligned} \quad (25)$$

Example 10.2. Consider interfaces $I_{00} := (\{x\}, \{y\}, x = 0 \rightarrow y = 0)$ and $I_{01} := (\{x\}, \{y\}, x = 0 \rightarrow y = 1)$. I_{00} and I_{01} are not shared-refinable because there is no way to satisfy $y = 0 \wedge y = 1$ when $x = 0$.

For finite-state interfaces, shared refinement is computable. Let $M_i = (X, Y, L_i, \ell_{0,i}, C_i, T_i)$ be finite-state automata representing I_i , for $i = 1, 2$, respectively. Suppose I_1, I_2 are shared-refinable. Then, $I_1 \sqcap I_2$ can be represented as the automaton




Fig. 12. Two interfaces that are not shared-refinable.

$M := (X, Y, L_1 \times L_2 \cup L_1 \cup L_2, (\ell_{0,1}, \ell_{0,2}), C, T)$, where C and T are defined as follows (guard g_{both} is defined as in (20)):

$$C(\ell) := \begin{cases} (\text{in}(C_1(\ell_1)) \vee \text{in}(C_2(\ell_2))) \wedge (\text{in}(C_1(\ell_1)) \rightarrow C_1(\ell_1) \wedge (\text{in}(C_2(\ell_2)) \rightarrow C_2(\ell_2))), & \text{if } \ell = (\ell_1, \ell_2) \in L_1 \times L_2 \\ C_1(\ell), & \text{if } \ell \in L_1 \\ C_2(\ell), & \text{if } \ell \in L_2 \end{cases} \quad (26)$$

$$T := \{((\ell_1, \ell_2), g_{both} \wedge g_1 \wedge g_2, (\ell'_1, \ell'_2)) \mid (\ell_i, g_i, \ell'_i) \in T_i, \text{ for } i = 1, 2\} \\ \cup \{((\ell_1, \ell_2), \neg C_2(\ell_2) \wedge g_1, \ell'_1) \mid (\ell_1, g_1, \ell'_1) \in T_1\} \cup T_1 \\ \cup \{((\ell_1, \ell_2), \neg C_1(\ell_1) \wedge g_2, \ell'_2) \mid (\ell_2, g_2, \ell'_2) \in T_2\} \cup T_2.$$

As long as the contracts of both M_1 and M_2 are satisfied, M behaves as a synchronous product. If the contract of one automaton is violated, then M continues with the other.

LEMMA 10.3. *If I and I' are shared-refinable interfaces, then*

$$f(I) \cap f(I') \subseteq f(I \sqcap I') \subseteq f(I) \cup f(I').$$

LEMMA 10.4. *Let I and I' be shared-refinable interfaces such that $I = (X, Y, f)$, $I' = (X, Y, f')$ and $I \sqcap I' = (X, Y, f_{\sqcap})$. Then for all $s \in f \cap f'$:*

$$\text{in}(f_{\sqcap}(s)) \equiv \text{in}(f(s)) \vee \text{in}(f'(s)).$$

THEOREM 10.5 (GREATEST LOWER BOUND). *If I and I' are shared-refinable interfaces, then $(I \sqcap I') \sqsubseteq I$, $(I \sqcap I') \sqsubseteq I'$, and for any interface I'' such that $I'' \sqsubseteq I$ and $I'' \sqsubseteq I'$, we have $I'' \sqsubseteq (I \sqcap I')$.*

Shared-refinability is a sufficient, but not necessary condition to existence of an interface I'' that refines both I and I' . The following example illustrates this fact.

Example 10.6. Consider interfaces I and I' shown in Figure 12. They have a single output variable y , and no inputs. I and I' are not shared-refinable. Indeed, $y = 1$ is initially possible in both interfaces, but after that, I requires $y = 0$ whereas I' requires $y = 1$, and there is no way of satisfying both. Nevertheless, an interface I'' exists that refines both I and I' : I'' is the stateless interface with contract $y = 0$.

THEOREM 10.7. *If I and I' are shared-refinable interfaces and both are well-formed, then $I \sqcap I'$ is well-formed.*

It is useful to consider the dual operator to \sqcap , that we call *shared abstraction* and denote \sqcup . Contrary to \sqcap , \sqcup is always defined, provided the interfaces have the same input and output variables.

Definition 10.8 (Shared abstraction). Two interfaces $I = (X, Y, f)$ and $I' = (X', Y', f')$ are *shared-abstractable* if $X = X'$ and $Y = Y'$. In that case, the *shared abstraction* of I and I' , denoted $I \sqcup I'$, is the interface:

$$I \sqcup I' := (X, Y, f_{\sqcup})$$

$$f_{\sqcup}(s) := \begin{cases} \text{in}(f(s)) \wedge \text{in}(f'(s)) \wedge (f(s) \vee f'(s)) & \text{if } s \in f \cap f' \\ f(s) & \text{if } s \in f \setminus f' \\ f'(s) & \text{if } s \in f' \setminus f \end{cases} \quad (27)$$

Notice that it suffices to define $f_{\sqcup}(s)$ for $s \in f \cup f'$. Indeed, the above definition inductively implies $f_{\sqcup} \subseteq f \cup f'$.

LEMMA 10.9. *If I and I' are shared-abstractable interfaces, then*

$$f(I) \cap f(I') \subseteq f(I \sqcup I') \subseteq f(I) \cup f(I').$$

For finite-state interfaces, shared abstraction is computable. Let $M_i = (X, Y, L_i, \ell_{0,i}, C_i, T_i)$ be finite-state automata representing I_i , for $i = 1, 2$, respectively. Suppose I_1, I_2 are shared-abstractable. Then, $I_1 \sqcup I_2$ can be represented as the automaton $M := (X, Y, L_1 \times L_2 \cup L_1 \cup L_2, (\ell_{0,1}, \ell_{0,2}), C, T)$, where C and T are defined as follows (guard g_{both} is defined as in (20)):

$$C(\ell) := \begin{cases} \text{in}(C_1(\ell_1)) \wedge \text{in}(C_2(\ell_2)) \wedge (L_1(\ell_1) \vee C_2(\ell_2)), & \text{if } \ell = (\ell_1, \ell_2) \in L_1 \times L_2 \\ C_1(\ell), & \text{if } \ell \in L_1 \\ C_2(\ell), & \text{if } \ell \in L_2 \end{cases} \quad (28)$$

$$T := \{((\ell_1, \ell_2), g_{\text{both}} \wedge g_1 \wedge g_2, (\ell'_1, \ell'_2)) \mid (\ell_i, g_i, \ell'_i) \in T_i, \text{ for } i = 1, 2\}$$

$$\cup \{((\ell_1, \ell_2), \text{in}(C_1(\ell_1)) \wedge \text{in}(C_2(\ell_2)) \wedge \neg C_2(\ell_2) \wedge g_1, \ell'_1) \mid (\ell_1, g_1, \ell'_1) \in T_1\} \cup T_1$$

$$\cup \{((\ell_1, \ell_2), \text{in}(C_1(\ell_1)) \wedge \text{in}(C_2(\ell_2)) \wedge \neg C_1(\ell_1) \wedge g_2, \ell'_2) \mid (\ell_2, g_2, \ell'_2) \in T_2\} \cup T_2.$$

Like the automaton for $I \sqcap I'$, M behaves as the synchronous product of M_1 and M_2 , as long as the contracts of both are satisfied. When the contract of one is violated, then M continues with the other.

THEOREM 10.10 (LEAST UPPER BOUND). *If I and I' are shared-abstractable interfaces, then $I \sqsubseteq (I \sqcup I')$, $I' \sqsubseteq (I \sqcup I')$, and for any interface I'' such that $I \sqsubseteq I''$ and $I' \sqsubseteq I''$, we have $(I \sqcup I') \sqsubseteq I''$.*

Notice that, even when I, I' are both well-formed, $I \sqcup I'$ may be non-well-formed, or even non-well-formable. This occurs, for instance, when I and I' are stateless with contracts ϕ and ϕ' such that $\text{in}(\phi) \wedge \text{in}(\phi')$ is false. This does not contradict Theorem 10.10 since false is refined by any contract, as observed earlier.

11. THE INPUT-COMPLETE CASE

Input-complete interfaces do not restrict the set of input values, although they may provide no guarantees when the input values are illegal. Although input-complete interfaces are a special case of general interfaces, it is instructive to study them separately for two reasons: first, input-completeness makes things much simpler, thus easier to understand and implement; second, some interesting properties hold for input-complete interfaces but not in general.

THEOREM 11.1. *Every well-formed Moore interface is input-complete.*

Note that source interfaces are Moore by definition, therefore every well-formed source interface is also input-complete.

THEOREM 11.2. *Every input-complete interface is well-formed.*

Every interface I can be transformed into an input-complete interface $\text{IC}(I)$. The illegal inputs of I become legal in $\text{IC}(I)$, but $\text{IC}(I)$ guarantees nothing about the value of the outputs when given such inputs. This transformation idea is well-known, for instance, it is called *chaotic closure* in Broy and Stølen [2001].

Definition 11.3 (Input-completion). Consider an interface $I = (X, Y, f)$. The *input-completion* of I , denoted $\text{IC}(I)$, is the interface $\text{IC}(I) := (X, Y, f_{ic})$, where $f_{ic}(s) := f(s) \vee \neg \text{in}(f(s))$, for all $s \in \mathcal{A}(X \cup Y)^*$.

THEOREM 11.4 (INPUT-COMPLETION REFINES ORIGINAL). *If I is an interface, then:*

- (1) $\text{IC}(I)$ is an input-complete interface.
- (2) $\text{IC}(I) \sqsubseteq I$.

Theorems 11.4 and 9.14 imply that for any environment E , if $I \models E$ then $\text{IC}(I) \models E$. The converse does not hold in general (see Examples 5.6 and 8.2, and observe that I_2 is the input-complete version of I_1).

Composition by connection reduces to conjunction of contracts for input-complete interfaces, and preserves input-completeness.

THEOREM 11.5. *Let $I_i = (X_i, Y_i, f_i)$, $i = 1, 2$, be disjoint input-complete interfaces, and let θ be a connection between I_1, I_2 . Then the contract f of the composite interface $\theta(I_1, I_2)$ is such that for all $s \in \mathcal{A}(X_{\theta(I_1, I_2)} \cup Y_{\theta(I_1, I_2)})^*$*

$$f(s) \equiv f_1(s) \wedge f_2(s) \wedge \rho_\theta.$$

Moreover, $\theta(I_1, I_2)$ is input-complete.

Input-complete interfaces alone do not help in avoiding problems with arbitrary feedback compositions: indeed, in the example given in the introduction both interfaces I_{true} and $I_{y \neq x}$ are input-complete.⁴ This means that in order to add a feedback connection (y, x) in an input-complete interface, we must still ensure that this interface is Moore w.r.t. input x . In that case, feedback preserves input-completeness.

THEOREM 11.6. *Let $I = (X, Y, f)$ be an input-complete interface which is also Moore with respect to some $x \in X$. Let $\kappa = (y, x)$ be a feedback connection on I . Then, $\kappa(I)$ is input-complete.*

THEOREM 11.7. *Let $I = (X, Y, f)$ be an input-complete interface and let $y \in Y$, such that f is independent from y . Then $\text{hide}(y, I)$ is input-complete.*

Theorem 11.8 follows directly from Definitions 9.1 and 5.3.

THEOREM 11.8 REFINEMENT FOR INPUT-COMplete INTERFACES. *Let I and I' be input-complete interfaces. Then $I' \sqsubseteq I$ iff $f(I') \sqsubseteq f(I)$.*

For input-complete interfaces, the shared-refinability condition, that is, Condition (24), simplifies to

$$\forall X : \exists Y : f(s) \wedge f'(s).$$

⁴It is not surprising that input-complete interfaces alone cannot solve the problems with arbitrary feedback compositions, since these are general problems of causality, not particular to interfaces.

Clearly, this condition does *not* always hold. Indeed, the interfaces of Example 10.2 are not shared-refinable, even though they are input-complete. For shared-refinable input-complete interfaces, shared refinement reduces to intersection. Dually, for shared-abstractable input-complete interfaces, shared abstraction reduces to union.

Theorem 11.9 follows directly from Definitions 10.1, 10.8 and 5.3.

THEOREM 11.9. *Let I and I' be input-complete interfaces.*

- (1) *If I and I' are shared-refinable, then $f(I \sqcap I') = f(I) \cap f(I')$.*
- (2) *If I and I' are shared-abstractable, then $f(I \sqcup I') = f(I) \cup f(I')$.*

12. THE DETERMINISTIC CASE

Deterministic interfaces produce a unique output for each legal input. As in the case of input-complete interfaces, it is instructive to study this sub-class of deterministic interfaces because the theory becomes simpler. Moreover, there is an interesting duality between the deterministic and input-complete case.

To begin, note that sink interfaces are by definition deterministic.

THEOREM 12.1. *All sink interfaces are deterministic.*

Composition by connection reduces to composition of relations when the source interface is deterministic.

THEOREM 12.2. *Consider two disjoint interfaces, $I_i = (X_i, Y_i, f_i)$, $i = 1, 2$, and a connection θ between I_1, I_2 . Let $\theta(I_1, I_2) = (X, Y, f)$. If I_1 is deterministic, then $f(s) \equiv f_1(s_1) \wedge f_2(s_2) \wedge \rho_\theta$ for all states s .*

THEOREM 12.3 (HIDING PRESERVES DETERMINISM). *Let $I = (X, Y, f)$ be a deterministic interface and let $y \in Y$, such that f is independent from y . Then $\text{hide}(y, I)$ is deterministic.*

THEOREM 12.4 (REFINEMENT FOR DETERMINISTIC INTERFACES). *Let I and I' be deterministic interfaces. Then $I' \sqsubseteq I$ iff $f(I') \supseteq f(I)$.*

A corollary of Theorems 11.8 and 12.4 is that refinement for input-complete and deterministic interfaces is equality.

For deterministic interfaces, the shared-refinability condition, that is, Condition (24), simplifies to

$$\forall X, Y : (\text{in}(f(s)) \wedge \text{in}(f'(s))) \rightarrow (f(s) \wedge f'(s))$$

Again, this condition does not always hold. For shared-refinable deterministic interfaces, shared refinement reduces to union. Dually, for shared-abstractable deterministic interfaces, shared abstraction reduces to intersection.

THEOREM 12.5. *Let I and I' be deterministic interfaces.*

- (1) *If I and I' are shared-refinable, then $f(I \sqcap I') = f(I) \cup f(I')$.*
- (2) *If I and I' are shared-abstractable, then $f(I \sqcup I') = f(I) \cap f(I')$.*

Notice that Theorems 12.4 and 12.5 are duals of Theorems 11.8 and 11.9.

13. APPLICATION: NONDEFENSIVE HARDWARE DESIGN

The theory developed in the previous sections is directly applicable to the domain of synchronous systems, which covers a broad class of applications, both in software and hardware. In particular, it applies to the class of applications captured in synchronous embedded software environments, as mentioned in the introduction. For instance, it




Fig. 13. Connecting a producer and a consumer.




Fig. 14. Stateful interfaces for a producer and a consumer.

can be used as a behavioral type theory for Simulink and other related models, in the spirit of Roy and Shankar [2010].

Synchronous hardware is another important application domain for our work. To illustrate this, we consider *nondefensive hardware design*, which is an application of Meyer’s ideas of *nondefensive programming* in the HW setting. To paraphrase Meyer, defensive programming consists in making SW modules input-complete, to guard against all possible inputs, including undesirable inputs that should not arise in principle [Meyer 1992]. Meyer argues that this is bad SW design practice, and we agree. Meyer proposes design-by-contract as an alternative. In a HW setting, the same defensive design practice is often encountered. Important benefits are to be obtained by abandoning this practice and by following the design-by-contract paradigm instead, to which our theory subscribes. We illustrate these points through an example.

Consider two HW components, Prod and Cons, having the input and output variables shown in Figure 13. Prod models a producer and Cons a consumer. Suppose that Cons requires that, once data starts being delivered at its input (i.e., once $valid_{in}$ becomes true), data continues to be delivered for 8 consecutive clock cycles. This is a typical requirement in HW “IP” (“intellectual property”) blocks that perform signal processing [Ravindran and Yang 2010].

We would like to connect Prod and Cons directly, as shown to the left of Figure 13. If we have no knowledge about Prod, however, we cannot do that, because Prod may produce data only intermittently, in which case the requirement of Cons is violated. Instead, we can insert a third component, Buff, to act as a mediator, as shown to the right of Figure 13. Buff acts as a temporary buffer that stores 8 values produced by Prod, and once 8 values become available, it signals and delivers them to Cons. The implementation details of Buff are not needed in this discussion. What is important is that Buff is an extra component that results in additional cost, both in terms of circuit size and performance (Cons must wait for Buff to accumulate 8 values before it starts processing them). We would like to avoid this cost. We can do this if we know that Prod conforms to the requirements of Cons: namely, that once Prod starts outputting data (i.e., once it sets $valid_{out}$ to true) it will continue to do so for 8 consecutive cycles. In that case, the direct connection of Prod to Cons is valid, and Buff becomes redundant.

The given situations can all be formally captured in our framework. Stateful interfaces can be used to model Prod and Cons, as shown in Figure 14. Interface Prod1 models a producer for which we have no knowledge, as in the first scenario described above. Interface Prod2 captures a different scenario where the producer is guaranteed

to produce 16 consecutive outputs once it starts producing data. Prod2 is captured as an automaton extended with an integer counter i ranging between 1 and 15. Since the domain of i is finite, Prod2 is a finite-state interface.⁵

The interface Cons for the consumer is also shown in Figure 14. The structure of Cons is similar to that of Prod2. Cons requires that $valid_{in}$ remains true for 8 consecutive rounds once it has been set to true. At the end of this period, the output $ready$ of Cons is set to true in order to signal that a batch of 8 consecutive inputs have been processed. Typically, Cons would also produce a value, but data values are completely abstracted in these interfaces. This results in simpler interfaces (with only a few states each), that can still be quite useful as this example illustrates.

Having these interfaces, we can formally state the fact that the unknown producer cannot be directly connected to our consumer. This is formalized by the fact that Prod1 and Cons are incompatible, that is, their serial composition is not well-formed (it is not well-formable either, since Prod1 has no inputs). On the other hand, we can formally state that Prod2 and Cons are compatible, therefore, an intermediate buffer is redundant in this case.

Note that the standard synchronous parallel composition of automata Prod1 and Cons does not reveal their incompatibility, since the conjunction of contracts true of Prod1 and $valid_{in}$ of Cons at its rightmost state, results in a satisfiable contract for the product state. On the other hand, a “demonic” interpretation of the nondeterminism of contract true of Prod1 reveals the error. In this simple example, where Prod1 has no inputs, this demonic interpretation can be easily captured by transforming Cons to an automaton Cons’ with an additional error location. This is similar to the error-completion transformation discussed in Section 2.4. Cons’ moves to the error location when an illegal input is received, that is, when $valid_{in}$ becomes false before 8 consecutive rounds have elapsed. Then, compatibility of Prod1 and Cons can be stated as a simple safety property on the standard parallel composition of Prod1 and Cons’, namely, that the error location of Cons’ is unreachable. This can be checked using a standard finite-state model-checker. In the general case, where Prod1 has inputs, compatibility cannot be stated as reachability and controller-synthesis algorithms must be used instead.

14. CONCLUSION AND PERSPECTIVES

We have proposed an interface theory that allows to reason formally about components and offers guarantees of substitutability. The framework we propose is general, and can be applied to a wide spectrum of cases, in particular within the synchronous model of computation. We are currently implementing our theory on the open-source Ptolemy software, and experimenting with different kinds of applications.

One major avenue for future work is to examine the current limitations on feedback compositions. Requiring feedback loops to contain Moore interfaces that “break” potential causality cycles is arguably a reasonable restriction in practice. After all, arbitrary feedback loops in synchronous models generally result in ambiguous semantics [Berry 1999; Malik 1994]. In many languages and tools these problems are avoided by making restrictions similar to (and often stricter than) ours. For example, Simulink and SCADE generally require a unit-delay to be present in every feedback loop. Similar restrictions are used in the synchronous language Lustre [Caspi et al. 1987].

Still, it would be interesting to study to what extent the current restrictions can be weakened. One possibility could be to refine the definition of Moore interfaces to

⁵Note that i is initialized to 1 and not to 0 when Prod2 switches from the initial location (with contract true) to the location with contract $valid_{out}$. This is because at this point one round where $valid_{out}$ was true already elapsed, namely, the round that triggered this transition when the automaton was at the initial location.

include dependencies between specific pairs of input and output variables. For example, this would allow one to express the fact that in the parallel composition of $(\{x_1\}, \{y_1\}, x_1 = y_1)$ and $(\{x_2\}, \{y_2\}, x_2 = y_2)$, y_1 does not depend on x_2 and y_2 does not depend on x_1 (and therefore one of the feedbacks (y_1, x_2) or (y_2, x_1) can be allowed). Such an extension could perhaps be achieved by combining our relational interfaces with the *causality interfaces* of Zhou and Lee [2008], input-output dependency information such as that used in reactive modules [Alur and Henzinger 1999], or the coarser *profiles* of Lubliner and Tripakis [2008]. A more general solution could involve studying fixpoints in a relational context, as is done, for instance, in Desharnais and Möller [2005].

In the current version of our theory contracts are prefix-closed sets, and therefore cannot express liveness properties. For instance, in the example of the buffer that may fail (Figure 9), we cannot express the requirement that if writes are attempted infinitely often then they must eventually succeed. In the future we plan to study extensions of the theory to handle liveness properties. It is worth noting, however, that the current theory already avoids the problem of trivial implementations that achieve the specification by “doing nothing.” An interface that “does nothing” is false, but false refines no other interface but itself. More generally, if an interface I is well-formed, then any refinement of I is well-formable, which means it can be executed forever without deadlocks.

Other directions of future work include examining canonical/minimal finite-state interfaces, as well as how nondeterministic automata can be used as representations of interfaces.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

ACKNOWLEDGMENTS

We are grateful to Jan Reineke, for helpful discussions on environments and error-complete interfaces, to Marc Geilen, who observed that the deterministic case is a dual of the input-complete case, and to Kaushik Ravindran and Guang Yang, for motivating the buffer examples. We would also like to thank Manfred Broy, Albert Benveniste, Rupak Majumdar and Slobodan Matic for their valuable feedback. We finally would like to thank the anonymous reviewers for their careful reading and comments that have greatly helped to improve this article.

REFERENCES

- ABADI, M. AND LAMPORT, L. 1995. Conjoining specifications. *ACM Trans. Program. Lang. Syst.* 17, 3, 507–535.
- ABRIAL, J.-R. 1996. *The B-book: Assigning Programs to Meanings*. Cambridge University Press, New York, NY.
- ALUR, R. AND HENZINGER, T. 1999. Reactive modules. *Form. Meth. Syst. Des.* 15, 7–48.
- ALUR, R., HENZINGER, T., KUPFERMAN, O., AND VARDI, M. 1998. Alternating refinement relations. In *Proceedings of the International Conference on Concurrency Theory (CONCUR'98)*. Lecture Notes in Computer Science, vol. 1466. Springer.
- BACK, R.-J. AND WRIGHT, J. 1998. *Refinement Calculus*. Springer.
- BARRINGER, H., KUIPER, R., AND PNUELI, A. 1984. Now you may compose temporal logic specifications. In *Proceedings of the 16th ACM Symposium on Theory of Computing (STOC'84)*. ACM, New York, NY, 51–63.
- BENVENISTE, A., CASPI, P., EDWARDS, S., HALBWACHS, N., LE GUERNIC, P., AND DE SIMONE, R. 2003. The synchronous languages 12 years later. *Proc. IEEE* 91, 1, 64–83.
- BENVENISTE, A., CAILLAUD, B., FERRARI, A., MANGERUCA, L., PASSERONE, R., AND SOFRONIS, C. 2008. Multiple viewpoint contract-based specification and design. In *Proceedings of the 6th International Symposium on Formal Methods for Components and Objects (FMCO'07)*. Springer, 200–225.

- BERRY, G. 1999. The constructive semantics of Pure Esterel. <http://www-sop.inria.fr/esterel.org/>.
- BROY, M. 1997. Compositional refinement of interactive systems. *J. ACM* 44, 6, 850–891.
- BROY, M. AND STØLEN, K. 2001. *Specification and Development of Interactive Systems: Focus on Streams, Interfaces, and Refinement*. Springer.
- CASPI, P., PILAUD, D., HALBWACHS, N., AND PLAICE, J. 1987. Lustre: A declarative language for programming synchronous systems. In *Proceedings of the 14th ACM Symposium on Principles of Programming Languages (POPL'87)*. ACM.
- CHAKRABARTI, A., DE ALFARO, L., HENZINGER, T., AND MANG, F. 2002. Synchronous and bidirectional component interfaces. In *Proceedings of the International Conference on Computer Aided Verification*. Lecture Notes in Computer Science vol. 2404, Springer, 414–427.
- CHEON, Y. AND LEAVENS, G. 1994. The Larch/Smalltalk interface specification language. *ACM Trans. Softw. Eng. Methodol.* 3, 3, 221–153.
- DE ALFARO, L. 2004. Game models for open systems. In *Verification: Theory and Practice*, N. Dershowitz Ed., Lecture Notes in Computer Science Series, vol. 2772, Springer, 192–213.
- DE ALFARO, L. AND HENZINGER, T. 2001a. Interface automata. In *Foundations of Software Engineering (FSE)*. ACM Press.
- DE ALFARO, L. AND HENZINGER, T. 2001b. Interface theories for component-based design. In *Proceedings of the International Workshop on Embedded Software (EMSOFT'01)*. Lecture Notes in Computer Science, vol. 2211, Springer.
- DESHARNAIS, J. AND MÖLLER, B. 2005. Least reflexive points of relations. *Higher Order Symbol. Comput.* 18, 1–2, 51–77.
- DHARA, K. AND LEAVENS, G. 1996. Forcing behavioral subtyping through specification inheritance. In *Proceedings of the 18th International Conference on Software Engineering (ICSE'96)*. IEEE Computer Society, 258–267.
- DIJKSTRA, E. 1972. Notes on structured programming. In *Structured Programming*, O. Dahl, E. Dijkstra, and C. Hoare Eds., Academic Press, London, UK, 1–82.
- DILL, D. 1987. *Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits*. MIT Press, Cambridge, MA.
- DOYEN, L., HENZINGER, T., JOBSTMANN, B., AND PETROV, T. 2008. Interface theories with component reuse. In *Proceedings of the 8th ACM & IEEE International Conference on Embedded Software*. 79–88.
- FLOYD, R. 1967. Assigning meanings to programs. In *Proceedings of the Symposium on Applied Mathematics*. American Mathematical Society, 19–32.
- FRAPPIER, M., MILI, A., AND DESHARNAIS, J. 1998. Unifying program construction and modification. *Logic J. IGPL* 6, 317–340.
- GRUMBERG, O. AND LONG, D. 1994. Model checking and modular verification. *ACM Trans. Program. Lang. Syst.* 16, 3, 843–871.
- GUTTAG, J. AND HORNING, J. 1993. *Larch: Languages and Tools for Formal Specification*. Springer.
- HEHNER, E. AND PARNAS, D. 1985. Technical correspondence. *Comm. ACM* 28, 5, 534–538.
- HENZINGER, T. AND SIFAKIS, J. 2007. The discipline of embedded systems design. *IEEE Computer* 40, 10, 32–40.
- HENZINGER, T., QADEER, S., AND RAJAMANI, S. 1998. You assume, we guarantee: Methodology and case studies. In *Proceedings of the International Conference on Computer-Aided Verification*. Lecture Notes in Computer Science, vol. 1427, Springer-Verlag.
- HOARE, C. A. R. 1969. An axiomatic basis for computer programming. *Comm. ACM* 12, 10, 576–580.
- HOARE, C. A. R. 1985. Programs are predicates. In *Proceedings of a Discussion Meeting of the Royal Society of London on Mathematical Logic and Programming Languages*. Prentice-Hall, Inc., Upper Saddle River, NJ, 141–155.
- JONES, C. B. 1983. Tentative steps toward a development method for interfering programs. *ACM Trans. Program. Lang. Syst.* 5, 4.
- JONSSON, B. 1994. Compositional specification and verification of distributed systems. *ACM Trans. Program. Lang. Syst.* 16, 2, 259–303.
- KAHL, W. 2003. Refinement and development of programs from relational specifications. *Electron. Notes Theor. Comput. Sci.* 44, 3, 51–93.
- LEAVENS, G. 1994. Inheritance of interface specifications. *SIGPLAN Notes* 29, 8, 129–138.
- LEAVENS, G. AND CHEON, Y. 2006. Design by contract with JML. <http://www.jmlspecs.org/jmldbc.pdf>.

- LEE, E. 2008. Cyber physical systems: Design challenges. Tech. rep. UCBIEECS-2008-8, EECS Department, University of California, Berkeley.
- LEE, E. AND SANGIOVANNI-VINCENTELLI, A. 1998. A unified framework for comparing models of computation. *IEEE Trans. Comput.-Aid. Des. Integr. Circ. Syst.* 17, 12, 1217–1229.
- LEE, E. AND XIONG, Y. 2001. System-level types for component-based design. In *Proceedings of the International Workshop on Embedded Software (EMSOFT'01)*. Springer, 237–253.
- LISKOV, B. 1979. Modular program construction using abstractions. In *Abstract Software Specifications*. Lecture Notes in Computer Science Series, vol. 86., Springer, 354–389.
- LISKOV, B. AND WING, J. 1994. A behavioral notion of subtyping. *ACM Trans. Program. Lang. Syst.* 16, 6, 1811–1841.
- LUBLINERMAN, R. AND TRIPAKIS, S. 2008. Modularity vs. Reusability: Code Generation from Synchronous Block Diagrams. In *Proceedings of the Conference and Exhibition on Design, Automation, and Test in Europe (DATE'08)*. ACM.
- LYNCH, N. AND TUTTLE, M. 1989. An introduction to input/output automata. *CWI Quart.* 2, 219–246.
- MALIK, S. 1994. Analysis of cyclic combinational circuits. *IEEE Trans. Comput.-Aid. Des.* 13, 7, 950–956.
- MCMILLAN, K. 1997. A compositional rule for hardware design refinement. In *Proceedings of the International Conference on Computer Aided Verification (CAV'97)*. Lecture Notes in Computer Science, vol. 1254, SpringerVerlag.
- MEYER, B. 1992. Applying “design by contract.” *Comput.* 25, 10, 40–51.
- MILLER, S., WHALEN, M., AND COFER, D. 2010. Software model checking takes off. *Comm. ACM* 53, 2, 58–64.
- MISRA, J. AND CHANDY, K. 1981. Proofs of networks of processes. *IEEE Trans. Softw. Engin.* 7, 4, 417–426.
- NELSON, G. 1989. A generalization of Dijkstra’s calculus. *ACM Trans. Program. Lang. Syst.* 11, 4, 517–561.
- NIERSTRASZ, O. 1993. Regular types for active objects. *SIGPLAN Notes* 28, 10, 1–15.
- PARNAS, D. 1983. A generalized control structure and its formal definition. *Comm. ACM* 26, 8, 572–581.
- PIERCE, B. 2002. *Types and Programming Languages*. MIT Press.
- RACLET, J.-B., BADOUEL, E., BENVENISTE, A., CAILLAUD, B., LEGAY, A., AND PASSERONE, R. 2010. A modal interface theory for component-based design. <http://www.irisa.fr/distribcom/benveniste/pub/Fundamenta2010.html>.
- RAVINDRAN, K. AND YANG, G. 2010. Personal communication.
- ROY, P. AND SHANKAR, N. 2010. An expressive type system for Simulink. In *Proceedings of the 2nd NASA Formal Methods Symposium (NFM'10)*. 149–160.
- SHANKAR, N. 1998. Lazy compositional verification. In *Compositionality: The Significant Difference*. Springer, 541–564.
- SPIVEY, J. M. 1989. *The Z Notation: A Reference Manual*. Prentice-Hall, Inc., Upper Saddle River, NJ.
- STARK, E. 1985. A proof technique for rely/guarantee properties. In *Proceedings of the 5th Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer-Verlag.
- TOURLAKIS, G. 2008. *Mathematical Logic*. Wiley.
- TRIPAKIS, S., LICKLY, B., HENZINGER, T., AND LEE, E. 2009a. On relational interfaces. Tech. rep. UCBIEECS-2009-60, EECS Department, University of California, Berkeley.
- TRIPAKIS, S., LICKLY, B., HENZINGER, T., AND LEE, E. 2009b. On relational interfaces. In *Proceedings of the 7th ACM International Conference on Embedded Software (EMSOFT'09)*. ACM, 67–76.
- WIRTH, N. 1971. Program development by stepwise refinement. *Comm. ACM* 14, 4, 221–227.
- ZHOU, Y. AND LEE, E. 2008. Causality interfaces for actor networks. *ACM Trans. Embed. Comput. Syst.* 7, 3, 1–35.

Received May 2010; revised November 2010; accepted March 2011