

The Refinement Calculus of Reactive Systems*

Viorel Preoteasa

Iulia Dragomir

Stavros Tripakis

October 12, 2017

Abstract

The Refinement Calculus of Reactive Systems (RCRS) is a compositional formal framework for modeling and reasoning about reactive systems. RCRS provides a language which allows to describe atomic components as symbolic transition systems or QLTL formulas, and composite components formed using three primitive composition operators: serial, parallel, and feedback. The semantics of the language is given in terms of monotonic property transformers, an extension to reactive systems of monotonic predicate transformers which have been used to give compositional semantics to sequential programs. RCRS allows to specify both safety and liveness properties. It also allows to model input-output systems which are both non-deterministic and non-input-receptive (i.e., which may reject some inputs at some points in time), and can thus be seen as a behavioral type system. RCRS provides a set of techniques for symbolic computer-aided reasoning, including compositional static analysis and verification. RCRS comes with an open-source implementation built on top of the Isabelle theorem prover.

Contents

1	Introduction	3
1.1	Novel contributions of this paper and relation to our prior work	4
2	Preliminaries	5
3	Language	6
3.1	An Algebra of Components	6
3.2	Symbolic Transition System Components	7
3.2.1	General STS Components	7
3.2.2	Variable Name Scope	9
3.2.3	Stateless STS Components	9
3.2.4	Deterministic STS Components	10
3.2.5	Stateless Deterministic STS Components	11
3.3	Quantified Linear Temporal Logic Components	11
3.3.1	QLTL	11
3.3.2	QLTL Components	13
3.4	Well Formed Components	13
4	Semantics	14
4.1	Monotonic Property Transformers	14
4.1.1	Some Commonly Used MPTs	15
4.1.2	Operators on MPTs: Function Composition, Product, and Fusion	16
4.1.3	Novel Operators Used in Semantical Definition of Feedback	17

*This work was partially supported by the Academy of Finland and the U.S. National Science Foundation (awards #1329759 and #1139138).

4.1.4	Refinement	18
4.2	Subclasses of MPTs	18
4.2.1	Relational MPTs	18
4.2.2	Guarded MPTs	19
4.2.3	Other subclasses and overview	20
4.3	Semantics of Components as MPTs	20
4.3.1	Example: Two Alternative Derivations of the Semantics of Diagram <code>Sum</code> of Fig. 2 . . .	21
4.3.2	Characterization of Legal Input Traces	22
4.3.3	Semantic Equivalence and Refinement for Components	22
4.3.4	Compositionality Properties	23
5	Symbolic Reasoning	23
5.1	Symbolic Transformation of STS Components to QLTL Components	23
5.2	Symbolic Transformations of Special Atomic Components to More General Atomic Components	25
5.3	Symbolic Computation of Serial Composition	26
5.3.1	Symbolic Serial Composition of Two QLTL Components	26
5.3.2	Symbolic Serial Composition of Two General STS Components	26
5.3.3	Symbolic Serial Composition of Two Stateless STS Components	26
5.3.4	Symbolic Serial Composition of Two Deterministic STS Components	27
5.3.5	Symbolic Serial Composition of Two Stateless Deterministic STS Components	27
5.3.6	Symbolic Serial Composition of Two Arbitrary Atomic Components	27
5.3.7	Correctness of Symbolic Serial Composition	27
5.4	Symbolic Computation of Parallel Composition	28
5.4.1	Symbolic Parallel Composition of Two QLTL Components	28
5.4.2	Symbolic Parallel Composition of Two General STS Components	28
5.4.3	Symbolic Parallel Composition of Two Stateless STS Components	29
5.4.4	Symbolic Parallel Composition of Two Deterministic STS Components	29
5.4.5	Symbolic Parallel Composition of Two Stateless Deterministic STS Components . . .	29
5.4.6	Symbolic Parallel Composition of Two Arbitrary Atomic Components	29
5.4.7	Correctness of Symbolic Parallel Composition	29
5.5	Symbolic Computation of Feedback Composition for Decomposable Deterministic STS Components	29
5.5.1	Decomposable Components	29
5.5.2	Symbolic Feedback of a Decomposable Deterministic STS Component	30
5.5.3	Symbolic Feedback of a Decomposable Stateless Deterministic STS Component . . .	30
5.5.4	Correctness of Symbolic Feedback Composition	30
5.6	Closure Properties of MPT Subclasses w.r.t. Composition Operators	31
5.7	Symbolic Simplification of Arbitrary Composite Components	31
5.7.1	Deterministic and Algebraic Loop Free Composite Components	31
5.7.2	Correctness of the Simplification Algorithm	33
5.8	Checking Validity and Compatibility	34
5.9	Checking Input-Receptiveness and Computing Legal Inputs Symbolically	35
5.10	Checking Refinement Symbolically	35
6	Toolset and Case Studies	37
7	Related Work	39
8	Conclusion	41

1 Introduction

This paper presents the *Refinement Calculus of Reactive Systems* (RCRS), a comprehensive framework for compositional modeling of and reasoning about reactive systems. RCRS originates from the precursor theory of *synchronous relational interfaces* [76, 77], and builds upon the classic *refinement calculus* [10]. A number of conference publications on RCRS exist [64, 31, 65, 63]. This paper collects some of these results and extends them in significant ways. The novel contributions of this paper and relation to our previous work are presented in §1.1.

The motivation for RCRS stems from the need for a compositional treatment of reactive systems. Generally speaking, compositionality is a divide-and-conquer principle. As systems grow in size, they grow in complexity. Therefore dealing with them in a monolithic manner becomes unmanageable. Compositionality comes to the rescue, and takes many forms. Many industrial-strength systems have employed for many years mechanisms for compositional modeling. An example is the Simulink tool from the Mathworks. Simulink is based on the widespread notation of hierarchical block diagrams. Such diagrams are both intuitive, and naturally compositional: a block can be refined into sub-blocks, sub-sub-blocks, and so on, creating hierarchical models of arbitrary depth. This allows the user to build large models (many thousands of blocks) while at the same time managing their complexity (at any level of the hierarchy, only a few blocks may be visible).

But Simulink’s compositionality has limitations, despite its hierarchical modeling approach. Even relatively simple problems, such as the problem of modular code generation (generating code for a block independently from context) require techniques not always available in standard code generators [50, 49]. Perhaps more serious, and more relevant in the context of this paper, is Simulink’s lack of formal semantics, and consequent lack of rigorous analysis techniques that can leverage the advances in the fields of computer-aided verification and programming languages.

RCRS provides a compositional formal semantics for Simulink in particular, and hierarchical block diagram notations in general, by building on well-established principles from the formal methods and programming language domains. In particular, RCRS relies on the notion of *refinement* (and its counterpart, *abstraction*) which are both fundamental in system design. Refinement is a binary relation between components, and ideally characterizes *substitutability*: the conditions under which some component can replace another component, without compromising the behavior of the overall system. RCRS refinement is *compositional* in the sense that it is preserved by composition: if A' refines A and B' refines B , then the composition of A' and B' refines the composition of A and B .

RCRS can be viewed as a refinement theory. It can also be viewed as a *behavioral type system*, similar to type systems for programming languages, but targeted to reactive systems. By *behavioral* we mean a type system that can capture not just data types of input and output ports of components (bool, int, etc.), but also complete specifications of the behavior of those components. As discussed more extensively in [77], such a behavioral type system has advantages over a full-blown verification system, as it is more lightweight: for instance, it allows type checking, which does not require the user to provide a formal specification of the properties that the model must satisfy: the model must simply *type-check*.

To achieve this, it is essential, as argued in [77], for the framework to be able to express *non-input-receptive* (also called *non-input-enabled* or *non-input-complete*) components, i.e., components that reject some input values. RCRS allows this. For example, a square-root component can be described in RCRS alternatively as: (1) a non-input-receptive component C_1 with input-output contract $x \geq 0 \wedge y = \sqrt{x}$ (where x, y are the input and output variables, respectively); or (2) an input-receptive component C_2 with contract $x \geq 0 \rightarrow y = \sqrt{x}$. Simple examples of type-checking in the RCRS context are the following: Connecting the non-input-receptive square-root component C_1 to a component which outputs $x = -1$ results in a type error (*incompatibility*). Connecting C_1 to a non-deterministic component which outputs an arbitrary value for x (this can be specified by the formula/contract *true*) also results in a type error.

RCRS allows both non-input-receptive and non-deterministic components. This combination results in a *game-theoretic* interpretation of composition operators, like in *interface theories* [22, 77]. Refinement also becomes game-theoretic, as in *alternating refinement* [7]. Game-theoretic composition can be used for an interesting form of type inference. For example, if we connect the non-input-receptive square-root component C_1 above to a non-deterministic component C_3 with input-output contract $x \geq u + 1$ (where x is the output

of C_3 , and u its input), and apply the (game-theoretic) serial composition of RCRS, we obtain the condition $u \geq -1$ on the external input of the overall composition. The constraint $u \geq -1$ represents the weakest condition on u which ensures compatibility of the connected components.

In a nutshell, RCRS consists of the following elements:

1. A modeling language (syntax), which allows to describe atomic components, and composite components, formed by connecting atomic components via a small number of primitive composition operators (serial, parallel, and feedback). The language is described in Section 3.
2. A formal semantics, presented in Section 4. Component semantics are defined in terms of *monotonic property transformers* (MPTs). MPTs are extensions of monotonic *predicate* transformers used in theories of programming languages, and in particular in refinement calculus [10]. Predicate transformers transform sets of *post-states* (states reached by the program after its computation) into sets of *pre-states* (states where the program begins). Property transformers transform sets of *output traces* (infinite sequences of input values) into sets of *input traces* (infinite sequences of output values). Using this semantics we can express both safety and liveness properties.
3. A set of symbolic reasoning techniques, described in Section 5. In particular, RCRS offers techniques to
 - compute the symbolic representation of a composite component from the symbolic representations of its sub-components;
 - simplify composite components into atomic components;
 - reduce checking refinement between two components to checking satisfiability of certain logical formulas;
 - reduce input-receptiveness and compatibility checks to satisfiability;
 - compute the legal inputs of a component symbolically.

We note that these techniques are for the most part *logic-agnostic*, in the sense that they do not depend on the particular logic used to represent components.

4. A toolset, described briefly in Section 6. The toolset consists mainly of:
 - a full implementation of the RCRS theory in the Isabelle proof assistant [58];
 - a translator of Simulink diagrams into RCRS code.

Our implementation is open-source and publicly available from <http://rcrs.cs.aalto.fi/>.

1.1 Novel contributions of this paper and relation to our prior work

RCRS was introduced in [64], which focuses on monotonic property transformers as a means to extend relational interfaces with liveness properties. [64] covers serial composition, but not parallel nor feedback. It also does not cover symbolic reasoning nor the RCRS toolset. Feedback is considered in [65], whose aim is in particular to study *instantaneous feedback* for non-deterministic and non-input-receptive systems. The study of instantaneous feedback is an interesting problem, but beyond the scope of the current paper. In this paper we consider non-instantaneous feedback, i.e., feedback for systems without same-step cyclic dependencies (no *algebraic loops*).

[31] presents the RCRS toolset, focusing on the translation of Simulink (and hierarchical block diagrams in general) into an algebra of components with three composition primitives, serial, parallel, and feedback, like RCRS. As it turns out, there is not a unique way to translate a graphical notation like Simulink into an algebraic formalism like RCRS. The problem of how exactly to do it and what are the trade-offs is an interesting one, but beyond the scope of the current paper. This problem is studied in depth in [31] which proposes three different translation strategies and evaluates their pros and cons. [31] leaves open the question

whether the results obtained by the different translations are equivalent. [61] settles this question, by proving that a class of translations, including the ones proposed in [31] are semantically equivalent for any input block diagram. [63] also concerns the RCRS toolset and discusses solutions to subtle typing problems that arise when translating Simulink diagrams into RCRS/Isabelle code.

In summary, the current paper does not cover the topics covered in [31, 61, 63, 65] and can be seen as a significantly revised and extended version of [64]. The main novel contributions with respect to [64] are the following: (1) a language of components (§3); (2) a revised MPT semantics (§4), including in particular novel operators for feedback (§4.1.3) and a classification of MPT subclasses (§4.2); (3) a new section of symbolic reasoning (§5).

2 Preliminaries

Sets, types. We use capital letters X, Y, Σ, \dots to denote types or sets, and small letters to denote elements of these types $x \in X, y \in Y$, etc. We denote by \mathbb{B} the type of Boolean values `true` and `false`. We use $\wedge, \vee, \rightarrow$, and \neg for the Boolean operations. The type of natural numbers is denoted by \mathbb{N} , while the type of real numbers is denoted by \mathbb{R} . The `Unit` type contains a single element `()`.

Cartesian product. For types X and Y , $X \times Y$ is the Cartesian product of X and Y , and if $x \in X$ and $y \in Y$, then (x, y) is a tuple from $X \times Y$. The empty Cartesian product is `Unit`. We assume that we have only flat products $X_1 \times \dots \times X_n$, and then we have

$$(X_1 \times \dots \times X_n) \times (Y_1 \times \dots \times Y_m) = X_1 \times \dots \times X_n \times Y_1 \times \dots \times Y_m$$

Functions. If X and Y are types, $X \rightarrow Y$ denotes the type of *functions* from X to Y . The function type constructor associates to the right (e.g., $X \rightarrow Y \rightarrow Z = X \rightarrow (Y \rightarrow Z)$) and the function interpretation associates to the left (e.g., $f(x)(y) = (f(x))(y)$). In order to construct functions we use lambda notation, e.g., $(\lambda x, y : x + y + 1) : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}$. Similarly, we can have tuples in the definition of functions, e.g., $(\lambda(x, y) : x + y + 2) : (\mathbb{N} \times \mathbb{N}) \rightarrow \mathbb{N}$. The composition of two functions $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, is a function denoted $g \circ f : X \rightarrow Z$, where $(g \circ f)(x) = g(f(x))$.

Predicates. A *predicate* is a function returning Boolean values, e.g., $p : X \rightarrow Y \rightarrow \mathbb{B}$, $p(x)(y) = (x = y)$. We define the *smallest predicate* $\perp : X \rightarrow \mathbb{B}$ where $\perp(x) = \text{false}$ for all $x \in X$. The *greatest predicate* is $\top : X \rightarrow \mathbb{B}$, with $\top(x) = \text{true}$ for all $x \in X$. We will often interpret predicates as sets. A predicate $p : X \rightarrow \mathbb{B}$ can be viewed as the set of all $x \in X$ such that $p(x) = \text{true}$. For example, viewing two predicates p, q as sets, we can write $p \subseteq q$, meaning that for all x , $p(x) \rightarrow q(x)$.

Relations. A *relation* is a predicate with at least two arguments, e.g., $r : X \rightarrow Y \rightarrow \mathbb{B}$. For such a relation r , we denote by $\text{in}(r) : X \rightarrow \mathbb{B}$ the predicate $\text{in}(r)(x) = (\exists y : r(x)(y))$. If the relation r has more than two arguments, then we define $\text{in}(r)$ similarly by quantifying over the last argument.

We extend point-wise all operations on Booleans to operations on predicates and relations. For example, if $r, r' : X \rightarrow Y \rightarrow \mathbb{B}$ are two relations, then $r \wedge r'$ and $r \vee r'$ are the relations given by $(r \wedge r')(x)(y) = r(x)(y) \wedge r'(x)(y)$ and $(r \vee r')(x)(y) = r(x)(y) \vee r'(x)(y)$. We also introduce the order on relations $r \subseteq r' = (\forall x, y : r(x)(y) \rightarrow r'(x)(y))$.

The composition of two relations $r : X \rightarrow Y \rightarrow \mathbb{B}$ and $r' : Y \rightarrow Z \rightarrow \mathbb{B}$ is a relation $(r \circ r') : X \rightarrow Z \rightarrow \mathbb{B}$, where $(r \circ r')(x)(z) = (\exists y : r(x)(y) \wedge r'(y)(z))$.

Infinite sequences. If Σ is a type, then $\Sigma^\omega = (\mathbb{N} \rightarrow \Sigma)$ is the set of all *infinite sequences* over Σ , also called *traces*. For a trace $\sigma \in \Sigma^\omega$, let $\sigma_i = \sigma(i)$ be the i -th element in the trace. Let $\sigma^i \in \Sigma^\omega$ denote the suffix of σ starting from the i -th step, i.e., $\sigma^i = \sigma_i \sigma_{i+1} \dots$. We often view a pair of traces $(\sigma, \sigma') \in \Sigma^\omega \times \Sigma'^\omega$ as being also a trace of pairs $(\lambda i : (\sigma_i, \sigma'_i)) \in (\Sigma \times \Sigma')^\omega$.

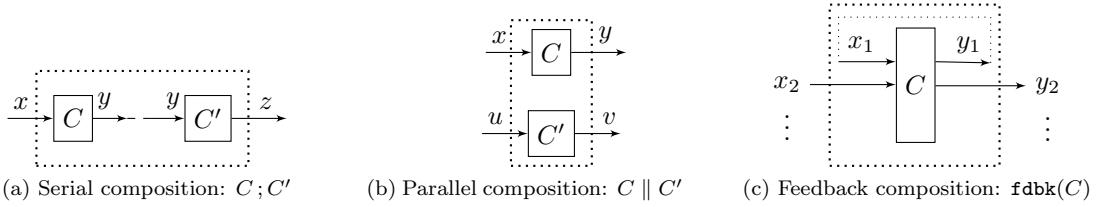


Figure 1: The three composition operators of RCRS.

Properties. A *property* is a predicate p over a set of infinite sequences. Formally, $p \in (\Sigma^\omega \rightarrow \mathbb{B})$. Just like any other predicate, a property can also be viewed as a set. In particular, a property can be viewed as a set of traces.

3 Language

3.1 An Algebra of Components

We model systems using a simple language of *components*. The grammar of the language is as follows:

```

component ::= atomic_component | composite_component
atomic_component ::= STS_component | QLTL_component
STS_component ::= GEN_STS_component | STATELESS_STS_component
                  | DET_STS_component | DET_STATELESS_STS_component
composite_component ::= component ; component | component || component | fdbk(component)

```

The elements of the above grammar are defined in the remaining of this section, where examples are also given to illustrate the language. In a nutshell, the language contains *atomic* components of two kinds: atomic components defined as *symbolic transition systems* (STS_components), and atomic components defined as QLTL formulas over input and output variables (QLTL_components).

STS components are split in four categories: general STS components, stateless STS components, deterministic STS components, and deterministic stateless STS components. Semantically, the general STS components subsume all the other more specialized STS components, but we introduce the specialized syntax because symbolic compositions of less general components become simpler, as we shall explain in the sequel (see §5).

Also, as it turns out, atomic components of our framework form a lattice, shown in Fig. 5, from the more specialized ones, namely, deterministic stateless STS components, to the more general ones, namely QLTL components. The full definition of this lattice will become apparent once we provide a symbolic transformation of STS to QLTL components, in §5.1.

Apart from atomic components, the language also allows to form *composite* components, by composing (atomic or other composite) components via three composition operators: serial $,$, parallel \parallel , and feedback \mathbf{fdbk} , as depicted in Fig. 1. The serial composition of two components C, C' is formed by connecting the output(s) of C to the input(s) of C' . Their parallel composition is formed by “stacking” the two components on top of each other without forming any new connections. The feedback of a component C is obtained by connecting the first output of C to its first input.

Our language is inspired by graphical notations such as Simulink, and hierarchical block diagrams in general. But our language is textual, not graphical. An interesting question is how to translate a graphical block diagram into a term in our algebra. We will not address this question here, as the issue is quite involved. We refer the reader to [31], which includes an extensive discussion on this topic. Suffice it to say here that there are generally many possible translations of a graphical diagram into a term in our algebra (or generally any algebra that contains primitive serial, parallel, and feedback composition operators). These

translations achieve different tradeoffs in terms of size, readability, computational properties, and so on. See [31] for details.

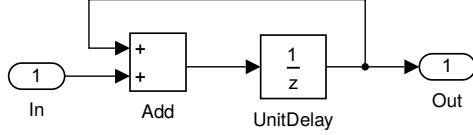


Figure 2: A Simulink diagram modeling the 1-step delayed sum of its input In . Each atomic block as well as the entire system can be formalized as STS components (see §3.2). The entire system can also be formalized as a composite component (see below).

As an example, consider the Simulink diagram shown in Fig. 2. This diagram can be represented in our language as a composite component `Sum` defined as

$$\text{Sum} = \text{fdbk}(\text{Add}; \text{UnitDelay}; \text{Split})$$

where `Add`, `UnitDelay`, and `Split` are atomic components (for a definition of these atomic components see §3.2). Here `Split` models the “fan-out” element in the Simulink diagram (black bullet) where the output wire of `UnitDelay` splits in two wires going to `Out` and back to `Add`.

3.2 Symbolic Transition System Components

We introduce all four categories of STS components and at the same time provide syntactic mappings from specialized STS components to general STS components.

3.2.1 General STS Components

A *general symbolic transition system component* (general STS component) is a transition system described symbolically, with Boolean expressions over input, output, state, and next state variables defining the initial states and the transition relation. When we say “Boolean expression” (here and in the definitions that follow) we mean an expression of type \mathbb{B} , in some arbitrary logic, not necessarily restricted to propositional logic. For example, if x is a variable of numerical type, then $x > 0$ is a Boolean expression. In the definition that follows, s' denotes the *primed*, or *next state* variable, corresponding to the *current state* variable s . Both can be vectors of variables. For example, if $s = (s_1, s_2)$ then $s' = (s'_1, s'_2)$. We assume that s' has the same type as s .

Definition 1 (STS component). *A (general) STS component is a tuple*

$$\text{sts}(x : \Sigma_x, y : \Sigma_y, s : \Sigma_s, \text{init_exp} : \mathbb{B}, \text{trs_exp} : \mathbb{B})$$

where x, y, s are input, output and state variables (or tuples of variables) of types $\Sigma_x, \Sigma_y, \Sigma_s$, respectively, init_exp is a Boolean expression on s (in some logic), and trs_exp is a Boolean expression on x, y, s, s' (in some logic).

Intuitively, an STS component is a non-deterministic system which for an infinite input sequence $\sigma_x \in \Sigma_x^\omega$ produces as output an infinite sequence $\sigma_y \in \Sigma_y^\omega$. The system starts non-deterministically at some state $\sigma_s(0)$ satisfying init_exp . Given first input $\sigma_x(0)$, the system non-deterministically computes output $\sigma_y(0)$ and next state $\sigma_s(1)$ such that trs_exp holds (if no such values exist, then the input $\sigma_x(0)$ is *illegal*, as discussed in more detail below). Next, it uses the following input $\sigma_x(1)$ and state $\sigma_s(1)$ to compute $\sigma_y(1)$ and $\sigma_s(2)$, and so on.

We will sometimes use the term *contract* to refer to the expression trs_exp . Indeed, trs_exp can be seen as specifying a contract between the component and its environment, in the following sense. At each step in the computation, the environment must provide input values that do not immediately violate the contract,

i.e., for which we can find values for the next state and output variables to satisfy trs_exp . Then, it is the responsibility of the component to find such values, otherwise it is the component’s “fault” if the contract is violated. This game-theoretic interpretation is similar in spirit with the classic refinement calculus for sequential programs [10].

We use Σ_x , Σ_y , Σ_s in the definition above to emphasize the types of the input, output and the state, and the fact that, when composing components, the types should match. However, in practice we often omit the types, unless they are required to unambiguously specify a component. Also note that the definition does not fix the logic used for the expressions init_exp and trs_exp . Indeed, our theory and results are independent from the choice of this logic. The choice of logic matters for algorithmic complexity and decidability. We will return to this point in §5. Finally, for simplicity, we often view the formulas init_exp and trs_exp as semantic objects, namely, as predicates. Adopting this view, init_exp becomes the predicate $\text{init} : \Sigma_s \rightarrow \mathbb{B}$, and trs_exp the predicate $\text{trs_exp} : \Sigma_s \rightarrow \Sigma_x \rightarrow \Sigma_s \rightarrow \Sigma_y \rightarrow \mathbb{B}$. Equivalently, trs_exp can be interpreted as a relation $\text{trs_exp} : (\Sigma_s \times \Sigma_x) \rightarrow (\Sigma_s \times \Sigma_y) \rightarrow \mathbb{B}$.

Examples. In the examples provided in this paper, we often specify systems that have tuples as input, state and output variables, in different equivalent ways. For example, we can introduce a general STS component with two inputs as $\text{sts}((n : \mathbb{N}, x : \mathbb{R}), s : \mathbb{R}, y : \mathbb{R}, s > 0, s' > s \wedge y + s = x^n)$, but also as $\text{sts}((n, x) : \mathbb{N} \times \mathbb{R}, s : \mathbb{R}, y : \mathbb{R}, s > 0, s' > s \wedge y + s = x^n)$, or $\text{sts}(z : \mathbb{N} \times \mathbb{R}, y : \mathbb{R}, s : \mathbb{R}, s > 0, s' > s \wedge y + s = \text{snd}(z)^{\text{fst}(z)})$, where fst and snd return the first and second elements of a pair.

Let us model a system that at every step i outputs the input received at previous step $i - 1$ (assume that the initial output value is 0). This corresponds to Simulink’s commonly used `UnitDelay` block, which is also modeled in the diagram of Fig. 2. This block can be represented by an STS component, where a state variable s is needed to store the input at moment i such that it can be used at the next step. We formally define this component as

$$\text{UnitDelay} = \text{sts}(x, y, s, s = 0, y = s \wedge s' = x).$$

We use first-order logic to define init_exp and trs_exp . Here init_exp is $s = 0$, which initializes the state variable s with the value 0. The trs_exp is $y = s \wedge s' = x$, that is at moment i the current state variable s which stores the input x received at moment $i - 1$ is outputted and its value is updated.

As another example, consider again the composite component `Sum` modeling the diagram of Fig. 2. `Sum` could also be defined as an atomic STS component:

$$\text{Sum} = \text{sts}(x, y, s, s = 0, y = s \wedge s' = s + x).$$

In §5 we will show how we can automatically and symbolically simplify composite component terms such as `fdbk(Add;UnitDelay;Split)`, to obtain syntactic representations of atomic components such as the one above.

These examples illustrate systems coming from Simulink models. However, our language is more general, and able to accomodate the description of other systems, such as state machines à la nuXmv [16], or input/output automata [51]. In fact, both `UnitDelay` and `Sum` are deterministic, so they could also be defined as deterministic STS components, as we will see below. Our language can capture non-deterministic systems easily. An example of a non-deterministic STS component is the following:

$$C = \text{sts}(x, y, s, s = 0, x + s \leq y).$$

For an input sequence $\sigma_x \in \mathbb{N}^\omega$, C outputs a non-deterministically chosen sequence σ_y such that the transition expression $x + s \leq y$ is satisfied. Since there is no formula in the transition expression tackling the next state variable, this is updated also non-deterministically with values from \mathbb{N} .

Our language can also capture *non-input-receptive* systems, that is, systems which disallow some input values as *illegal*. For instance, a component performing division, but disallowing division by zero, can be specified as follows:

$$\text{Div} = \text{sts}((x, y), z, (), \text{true}, y \neq 0 \wedge z = \frac{x}{y}).$$

Note that `Div` has an empty tuple of state variables, $s = ()$. Such components are called *stateless*, and are introduced next.

Even though RCRS is primarily a discrete-time framework, we have used it to model and verify continuous-time systems such as those modeled in Simulink (see Section 6). We do this by discretizing time using a time step parameter $\Delta t > 0$ and applying Euler numerical integration. Then, we can model Simulink’s *Integrator* block in RCRS as an STS component parameterized by Δt :

$$\text{Integrator}_{\Delta t} = \text{sts}(x, y, s, s = 0, y = s \wedge s' = s + x \cdot \Delta t)$$

More complex dynamical system blocks can be modeled in a similar fashion. For instance, Simulink’s *Transfer Fcn* block, with transfer function

$$\frac{s^2 + 2}{0.5s^2 + 2s + 1}$$

can be modeled in RCRS as the following STS component parameterized by Δt :

$$\begin{aligned} \text{TransferFcn}_{\Delta t} &= \text{sts}(x, y, (s_1, s_2), s_1 = 0 \wedge s_2 = 0, trs) \\ \text{where } trs &= (y = -8 \cdot s_1 + 2 \cdot x) \wedge \\ &\quad (s'_1 = s_1 + (-4 \cdot s_1 - 2 \cdot s_2 + x) \cdot \Delta t) \wedge \\ &\quad (s'_2 = s_2 + s_1 \cdot \Delta t) \end{aligned}$$

3.2.2 Variable Name Scope

We remark that variable names in the definition of atomic components are *local*. This holds for all atomic components in the language of RCRS (including STS and QLTL components, defined in the sequel). This means that if we replace a variable with another one in an atomic component, then we obtain a semantically equivalent component. For example, the two STS components below are equivalent (the semantical equivalence symbol \equiv will be defined formally in Def. 20, once we define the semantics):

$$\text{sts}((x, y), z, s, s > 0, z > x + y) \equiv \text{sts}((u, v), w, t, t > 0, w > u + v)$$

3.2.3 Stateless STS Components

A special STS component is one that has no state variables:

Definition 2 (Stateless STS component). *A stateless STS component is a tuple*

$$C = \text{stateless}(x : \Sigma_x, y : \Sigma_y, io_exp : \mathbb{B})$$

where x, y are the input and output variables, and io_exp is a Boolean expression on x and y . Stateless STS components are special cases of general STS components, as defined by the mapping `stateless2sts`:

$$\text{stateless2sts}(C) = \text{sts}(x, y, (), \text{true}, io_exp).$$

Examples. A trivial stateless STS component is the one that simply transfers its input to its output (i.e., a “wire”). We denote such a component by `Id`, and we formalize it as

$$\text{Id} = \text{stateless}(x, y, y = x).$$

Component `Add` from Fig. 2, which outputs the sum of its two inputs, can be modeled as a stateless STS component:

$$\text{Add} = \text{stateless}((x, y), z, z = x + y).$$

Component `Split` from Fig. 2 can also be modeled as a stateless STS component:

$$\text{Split} = \text{stateless}(x, (y, z), y = x \wedge z = x).$$

The `Div` component introduced above is stateless, and therefore can be also specified as follows:

$$\text{Div} = \text{stateless}((x, y), z, y \neq 0 \wedge z = \frac{x}{y}).$$

The above examples are not only stateless, but also deterministic. We introduce deterministic STS components next.

3.2.4 Deterministic STS Components

Deterministic STS components are those which, for given current state and input, have at most one output and next state. Syntactically, they are introduced as follows:

Definition 3 (Deterministic STS component). *A deterministic STS component is a tuple*

$$\text{det}(x : \Sigma_x, s : \Sigma_s, a : \Sigma_s, \text{inpt_exp} : \mathbb{B}, \text{next_exp} : \Sigma_s, \text{out_exp} : \Sigma_y)$$

where x, s are the input and state variables, $a \in \Sigma_s$ is the initial value of the state variable, inpt_exp is a Boolean expression on s and x defining the legal inputs, next_exp is an expression of type Σ_s on x and s defining the next state, and out_exp is an expression of type Σ_y on x and s defining the output. Deterministic STS components are special cases of general STS components, as defined by the mapping `det2sts`:

$$\text{det2sts}(C) = (x, y, s, s = a, \text{inpt_exp} \wedge s' = \text{next_exp} \wedge y = \text{out_exp})$$

where y is a new variable name (or tuple of new variable names) of type Σ_y .

Note that a deterministic STS component has a separate expression `inpt_exp` to define legal inputs. A separate such expression is not needed for general STS components, where the conditions for legal inputs are part of the expression `trs_exp`. For example, compare the definition of `Div` as a general STS above, and as a stateless deterministic STS below (see §3.2.5).

Examples. As mentioned above, all three components, `UnitDelay`, `Add`, and `Split`, from Fig. 2, are deterministic. They could therefore be specified in our language as deterministic STS components, instead of general STS components:

$$\begin{aligned} \text{UnitDelay} &= \text{det}(x, s, 0, \text{true}, x, s) \\ \text{Add} &= \text{det}((x, y), (), (), \text{true}, (), x + y) \\ \text{Split} &= \text{det}(x, (), (), \text{true}, (), (x, x)) \end{aligned}$$

The component `Sum` modeling the entire system is also deterministic, and could be defined as a deterministic STS component:

$$\text{Sum} = \text{det}(x, s, 0, \text{true}, s + x, s).$$

Note that these alternative specifications for each of those components, although syntactically distinct, will turn out to be semantically equivalent by definition, when we introduce the semantics of our language, in §4.

3.2.5 Stateless Deterministic STS Components

STS components which are both deterministic and stateless can be specified as follows:

Definition 4 (Stateless deterministic STS component). *A stateless deterministic STS component is a tuple*

$$C = \text{stateless_det}(x : \Sigma_x, \text{inpt_exp} : \mathbb{B}, \text{out_exp} : \Sigma_y)$$

where x is the input variable, inpt_exp is a Boolean expression on x defining the legal inputs, and out_exp is an expression of type Σ_y on x defining the output. Stateless deterministic STS components are special cases of both deterministic STS components, and of stateless STS components, as defined by the mappings

$$\begin{aligned} \text{stateless_det2det}(C) &= \text{det}(x, (), (), \text{inpt_exp}, (), \text{out_exp}) \\ \text{stateless_det2stateless}(C) &= \text{stateless}(x, y, \text{inpt_exp} \wedge y = \text{out_exp}) \end{aligned}$$

where y is a new variable name or a tuple of new variable names.

Examples. Many of the examples introduced above are both deterministic and stateless. They could be specified as follows:

$$\begin{aligned} \text{Id} &= \text{stateless_det}(x, \text{true}, x) \\ \text{Add} &= \text{stateless_det}((x, y), \text{true}, x + y) \\ \text{Split} &= \text{stateless_det}(x, \text{true}, (x, x)) \\ \text{Div} &= \text{stateless_det}((x, y), y \neq 0, \frac{x}{y}) \end{aligned}$$

3.3 Quantified Linear Temporal Logic Components

Although powerful, STS components have limitations. In particular, they cannot express *liveness* properties [4]. To remedy this, we introduce another type of components, based on Linear Temporal Logic (LTL) [60] and quantified propositional LTL (QPTL) [74, 45], which extends LTL with \exists and \forall quantifiers over propositional variables. In this paper we use *quantified first-order LTL* (which we abbreviate as QLTL). QLTL further extends QPTL with functional and relational symbols over arbitrary domains, quantification of variables over these domains, and a next operator applied to variables.¹ We need this expressive power in order to be able to handle general models (e.g., Simulink) which often use complex arithmetic formulas, and also to be able to translate STS components into semantically equivalent QLTL components (see §5.1).

3.3.1 QLTL

QLTL formulas are generated by the following grammar. We assume a set of constants and functional symbols ($0, 1, \dots, \text{true}, \text{false}, +, \dots$), a set of predicate symbols ($=, \leq, <, \dots$), and a set of variable names (x, y, z, \dots).

Definition 5 (Syntax of QLTL). *A QLTL formula φ is defined by the following grammar:*

$$\begin{array}{lll} \text{term} & ::= & x \mid y \mid \dots \mid & (\text{variable names}) \\ & & 0 \mid 1 \mid \dots \mid \text{true} \mid \dots \mid & (\text{constants}) \\ & & \text{term} + \text{term} \mid \dots \mid & (\text{functional symbol application}) \\ & & \circ \text{term} & (\text{next applied to a term}) \\ \varphi & ::= & \text{term} = \text{term} \mid (\text{term} \leq \text{term}) \mid \dots \mid & (\text{atomic QLTL formulas}) \\ & & \neg \varphi \mid & (\text{negation}) \\ & & \varphi \vee \psi \mid & (\text{disjunction}) \\ & & \varphi \mathbf{U} \psi \mid & (\text{until}) \\ & & \forall x : \varphi & (\text{forall}) \end{array}$$

¹ A logic similar to the one that we use here is presented in [44], however in [44] the next operator can be applied only once to variables, and the logic from [44] uses also past temporal operators.

As in standard first order logic, the *bounded variables* of a formula φ are the variables in scope of the universal quantifier \forall , and the *free variables* of φ are those that are not bounded. The logic connectives \wedge , \Rightarrow and \Leftrightarrow can be expressed with \neg and \vee . Quantification is over atomic variables. The existential quantifier \exists can be defined via the universal quantifier usually as $\neg\forall\neg$. The primitive temporal operators are *next for terms* (\circlearrowright) and *until* (\mathbf{U}). As is standard, QLTL formulas are evaluated over infinite traces, and $\varphi \mathbf{U} \psi$ intuitively means that φ continuously holds until some point in the trace where ψ holds.

Formally, we will define the relation $\sigma \models \varphi$ (σ satisfies φ) for a QLTL formula φ over free variables x, y, \dots , and an infinite sequence $\sigma \in \Sigma^\omega$, where $\Sigma = \Sigma_x \times \Sigma_y \times \dots$, and $\Sigma_x, \Sigma_y, \dots$ are the types (or *domains*) of variables x, y, \dots . As before we assume that σ can be written as a tuple of sequences $(\sigma_x, \sigma_y, \dots)$ where $\sigma_x \in \Sigma_x^\omega, \sigma_y \in \Sigma_y^\omega, \dots$. The semantics of a term t on variables x, y, \dots is a function from infinite sequences to infinite sequences $\langle\langle t \rangle\rangle : \Sigma^\omega \rightarrow \Sigma_t^\omega$, where $\Sigma = \Sigma_x \times \Sigma_y \times \dots$, and Σ_t is the type of t . When giving the semantics of terms and formulas we assume that constants, functional symbols, and predicate symbols have the standard semantics. For example, we assume that $+, \leq, \dots$ on numeric values have the semantics of standard arithmetic.

Definition 6 (Semantics of QLTL). *Let x be a variable, t, t' be terms, φ, ψ be QLTL formulas, P be a predicate symbol, f be a functional symbol, c be a constant, and $\sigma \in \Sigma^\omega$ be an infinite sequence.*

$$\begin{aligned}\langle\langle x \rangle\rangle(\sigma) &:= \sigma_x \\ \langle\langle c \rangle\rangle(\sigma) &:= (\lambda i : c) \\ \langle\langle f(t, t') \rangle\rangle(\sigma) &:= (\lambda i : f(\langle\langle t \rangle\rangle(\sigma)(i), \langle\langle t' \rangle\rangle(\sigma)(i))) \\ \langle\langle \circlearrowright t \rangle\rangle(\sigma) &:= \langle\langle t \rangle\rangle(\sigma^1) \\ \sigma \models P(t, t') &:= P(\langle\langle t \rangle\rangle(\sigma)(0), \langle\langle t' \rangle\rangle(\sigma)(0)) \\ \sigma \models \neg\varphi &:= \neg \sigma \models \varphi \\ \sigma \models \varphi \vee \psi &:= \sigma \models \varphi \vee \sigma \models \psi \\ \sigma \models \varphi \mathbf{U} \psi &:= (\exists n \geq 0 : (\forall 0 \leq i < n : \sigma^i \models \varphi) \wedge \sigma^n \models \psi) \\ \sigma \models (\forall x : \varphi) &:= (\forall \sigma_x \in \Sigma_x^\omega : (\sigma_x, \sigma) \models \varphi)\end{aligned}$$

Other temporal operators can be defined as follows. *Eventually* ($\mathbf{F} \varphi = \mathbf{true} \mathbf{U} \varphi$) states that φ must hold in some future step. *Always* ($\mathbf{G} \varphi = \neg\mathbf{F} \neg\varphi$) states that φ must hold at all steps. The next operator for formulas \mathbf{X} can be defined using the next operator for terms \circlearrowright . The formula $\mathbf{X}\varphi$ is obtained by replacing all occurrences of the free variables in φ by their next versions (i.e., x is replaced by $\circlearrowright x$, y by $\circlearrowright y$, etc.). For example the propositional LTL formula $\mathbf{X}(x \wedge \mathbf{X} y \Rightarrow \mathbf{G} z)$ can be expressed as

$$(\circlearrowright x = \mathbf{true} \wedge (\circlearrowright \circlearrowright y = \mathbf{true}) \Rightarrow \mathbf{G}(\circlearrowright z = \mathbf{true})).$$

We additionally introduce the operator: $\varphi \mathbf{L} \psi = \neg(\varphi \mathbf{U} \neg\psi)$. Intuitively, $\varphi \mathbf{L} \psi$ holds if whenever φ holds continuously up to some step $n - 1$, ψ must hold at step n .

Two QLTL formulas φ and ψ are semantically equivalent, denoted $\varphi \equiv \psi$, if

$$\forall \sigma : (\sigma \models \varphi) \Leftrightarrow (\sigma \models \psi).$$

Lemma 1. *Let φ be a QLTL formula. Then:*

1. $(\exists x : \mathbf{G} \varphi) \equiv \mathbf{G}(\exists x : \varphi)$ when φ does not contain temporal operators.
2. $\varphi \mathbf{L} \varphi \equiv \mathbf{G} \varphi$
3. $\mathbf{true} \mathbf{L} \varphi \equiv \mathbf{G} \varphi$
4. $\varphi \mathbf{L} \mathbf{true} \equiv \mathbf{true}$
5. $\varphi \mathbf{L} \mathbf{false} \equiv \mathbf{false}$
6. $\forall y : (\varphi \mathbf{L} \psi) \equiv (\exists y : \varphi) \mathbf{L} \psi$, when φ does not contain temporal operators and y is not free in ψ .

The proof of the above result, as well as most results that follow, is omitted. All omitted proofs have been mechanized in Isabelle, and are available as part of the public distribution of RCRS from rcrs.cs.aalto.fi.

Examples. Using QLTL we can express *safety*, as well as *liveness requirements*. Informally, a safety requirement expresses that something bad never happens. An example is the formula

$$\text{thermostat} = \mathbf{G}(180^\circ \leq t \wedge t \leq 220^\circ),$$

which states that the thermostat-controlled temperature t stays always between 180° and 220° .

A liveness requirement informally says that something good eventually happens. An example is the formula $\mathbf{F}(t > 200^\circ)$ stating that the temperature t is eventually over 200° .

A more complex example is a formula modeling an oven that starts increasing the temperature from an initial value of 20° until it reaches 180° , and then keeps it between 180° and 220° .

$$\text{oven} = (t = 20^\circ \wedge ((t < \bigcirc t \wedge t < 180) \mathbf{U} \text{thermostat})).$$

In this example the formula $t < \bigcirc t$ specifies that the temperature increases from some point to the next.

3.3.2 QLTL Components

A QLTL component is an atomic component where the input-output behavior is specified by a QLTL formula:

Definition 7 (QLTL component). A QLTL component is a tuple $\text{qltl}(x : \Sigma_x, y : \Sigma_y, \varphi)$, where x, y are input and output variables (or tuples of variables) of types Σ_x, Σ_y , and φ is a QLTL formula over x and y .

Intuitively a QLTL component $C = \text{qltl}(x, y, \varphi)$ represents a system that takes as input an infinite sequence $\sigma_x \in \Sigma_x^\omega$ and produces as output an infinite sequence $\sigma_y \in \Sigma_y^\omega$ such that $(\sigma_x, \sigma_y) \models \varphi$. If there is no σ_y such that $(\sigma_x, \sigma_y) \models \varphi$ is true, then input σ_x is illegal for C , i.e., C is not input-receptive. There could be many possible σ^y for a single σ_x , in which case the system is non-deterministic.

As a simple example, we can model the oven as a QLTL component with no input variables and the temperature as the only output variable:

$$\text{qltl}(\emptyset, t, \text{oven})$$

3.4 Well Formed Components

Not all of the terms generated by the grammar introduced in §3.1 are *well formed*. For example, two STS components C and C' can be composed in series only if the number of outputs of C matches the number of inputs of C' , and in addition the input types of C' are the same as the corresponding output types of C . Formally, for every component C we define below $\Sigma_{in}(C)$ - the *input type* of C , $\Sigma_{out}(C)$ - the *output type* of C , and $\text{wf}(C)$ - the *well-formedness* of C by induction on the structure of C .

$$\begin{aligned} \Sigma_{in}(\text{sts}(x : \Sigma_x, y : \Sigma_y, s : \Sigma_s, init, trs)) &= \Sigma_x \\ \Sigma_{in}(\text{stateless}(x : \Sigma_x, y : \Sigma_y, trs)) &= \Sigma_x \\ \Sigma_{in}(\text{det}(x : \Sigma_x, s : \Sigma_s, a, inpt, next, out : \Sigma_y)) &= \Sigma_x \\ \Sigma_{in}(\text{stateless_det}(x : \Sigma_x, inpt, out : \Sigma_y)) &= \Sigma_x \\ \Sigma_{in}(\text{qltl}(x : \Sigma_x, y : \Sigma_y, \varphi)) &= \Sigma_x \\ \Sigma_{in}(C ; C') &= \Sigma_{in}(C) \\ \Sigma_{in}(C \parallel C') &= \Sigma_{in}(C) \times \Sigma_{in}(C') \\ \Sigma_{in}(\text{fdbk}(C)) &= X_2 \times \dots \times X_n \text{ where } \Sigma_{in}(C) = X_1 \times \dots \times X_n, n > 0 \end{aligned}$$

$$\begin{aligned} \Sigma_{out}(\text{sts}(x : \Sigma_x, y : \Sigma_y, s : \Sigma_s, init, trs)) &= \Sigma_y \\ \Sigma_{out}(\text{stateless}(x : \Sigma_x, y : \Sigma_y, trs)) &= \Sigma_y \\ \Sigma_{out}(\text{det}(x : \Sigma_x, s : \Sigma_s, a, inpt, next, out : \Sigma_y)) &= \Sigma_y \\ \Sigma_{out}(\text{stateless_det}(x : \Sigma_x, inpt, out : \Sigma_y)) &= \Sigma_y \\ \Sigma_{out}(\text{qltl}(x : \Sigma_x, y : \Sigma_y, \varphi)) &= \Sigma_y \\ \Sigma_{out}(C ; C') &= \Sigma_{out}(C') \\ \Sigma_{out}(C \parallel C') &= \Sigma_{out}(C) \times \Sigma_{out}(C') \\ \Sigma_{out}(\text{fdbk}(C)) &= Y_2 \times \dots \times Y_n \text{ where } \Sigma_{out}(C) = Y_1 \times \dots \times Y_n, n > 0 \end{aligned}$$

$$\begin{aligned}
\text{wf}(\text{sts}(x, y, s, \text{init}, \text{trs})) &= \text{true} \\
\text{wf}(\text{stateless}(x, y, \text{trs})) &= \text{true} \\
\text{wf}(\text{det}(x, s, a, \text{inpt}, \text{next}, \text{out})) &= \text{true} \\
\text{wf}(\text{stateless_det}(x, \text{inpt}, \text{out})) &= \text{true} \\
\text{wf}(\text{qltl}(x, y, \varphi)) &= \text{true} \\
\text{wf}(C ; C') &= \text{wf}(C) \wedge \text{wf}(C') \wedge \Sigma_{\text{out}}(C) = \Sigma_{\text{in}}(C') \\
\text{wf}(C \parallel C') &= \text{wf}(C) \wedge \text{wf}(C') \\
\text{wf}(\text{fdbk}(C)) &= \text{wf}(C) \wedge \Sigma_{\text{in}}(C) = X \times X_1 \dots \times X_n \wedge \Sigma_{\text{out}}(C) = X \times Y_1 \dots \times Y_n
\end{aligned}$$

We note that the above well-formedness conditions are not restrictive. Components that do not have matching inputs and outputs can still be composed by adding appropriate *switching* components which reorder inputs, duplicate inputs, and so on. An example of such a component is the component **Split**, introduced earlier. As another example, consider the diagram in Fig. 3:

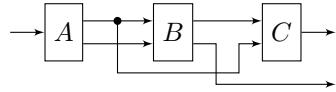


Figure 3: Another block diagram.

This diagram can be expressed in our language as the composite component:

$$A ; \text{Switch1} ; (B \parallel \text{Id}) ; \text{Switch2} ; (C \parallel \text{Id})$$

where

$$\begin{aligned}
\text{Switch1} &= \text{stateless_det}((x, y), \text{true}, (x, y, x)) \\
\text{Switch2} &= \text{stateless_det}((u, v, x), \text{true}, (u, x, v))
\end{aligned}$$

4 Semantics

In RCRS, the semantics of components is defined in terms of *monotonic property transformers* (MPTs). This is inspired by classical refinement calculus [10], where the semantics of sequential programs is defined in terms of monotonic *predicate* transformers [28]. Predicate transformers are functions that transform sets of *post-states* (states reached after the program executes) into sets of *pre-states* (states from which the program begins). Property transformers map sets of *output traces* (that a component produces) into sets of *input traces* (that a component consumes).

In this section we define MPTs formally, and introduce some basic operations on them, which are necessary for giving the semantics of components. The definitions of some of these operations (e.g., product and fusion) are simple extensions of the corresponding operations on predicate transformers [10, 9]. Other operations, in particular those related to feedback, are new (§4.1.3). The definition of component semantics is also new (§4.3).

4.1 Monotonic Property Transformers

A property transformer is a function $S : (\Sigma_y^\omega \rightarrow \text{Bool}) \rightarrow (\Sigma_x^\omega \rightarrow \text{Bool})$, where Σ_x, Σ_y are input and output types of the component in question. Note that x is the input and y is the output. A property transformer has a weakest precondition interpretation: it is applied to a set of output traces $Q \subseteq \Sigma_y^\omega$, and returns a set of input traces $P \subseteq \Sigma_x^\omega$, such that all traces in P are legal and, when fed to the component, are guaranteed to produce only traces in Q as output.

Interpreting properties as sets, monotonicity of property transformers simply means that these functions are monotonic with respect to set inclusion. That is, S is *monotonic* if for any two properties q, q' , if $q \subseteq q'$ then $S(q) \subseteq S(q')$.

For an MPT S we define its set of *legal input traces* as $\text{legal}(S) = S(\top)$, where \top is the greatest predicate extended to traces. Note that, because of monotonicity, and the fact that $q \subseteq \top$ holds for any property q , we have that $S(q) \subseteq \text{legal}(S)$ for all q . This justifies the definition of $\text{legal}(S)$ as a “maximal” set of input traces for which a system does not fail, assuming no restrictions on the post-condition. An MPT S is said to be *input-receptive* if $\text{legal}(S) = \top$.

4.1.1 Some Commonly Used MPTs

Definition 8 (Skip). *Skip is defined to be the MPT such that for all q , $\text{Skip}(q) = q$.*

Skip models the identity function, i.e., the system that accepts all input traces and simply transfers them unchanged to the output (this will become more clear when we express Skip in terms of assert or update transformers, below). Note that Skip is different from Id, defined above, although the two are strongly related: Id is a component, i.e., a syntactic object, while Skip is an MPT, i.e., a semantic object. As we shall see in §4.3, the semantics of Id is defined as Skip.

Definition 9 (Fail). *Fail is defined to be the MPT such that for all q , $\text{Fail}(q) = \perp$.*

Recall that \perp is the predicate that returns false for any input. Thus, viewed as a set, \perp is the empty set. Consequently, Fail can be seen to model a system which rejects all inputs, i.e., a system such that for any output property q , there are no input traces that can produce an output trace in q .

Definition 10 (Assert). *Let $p \in \Sigma^\omega \rightarrow \text{Bool}$ be a property. The assert property transformer $\{p\} : (\Sigma^\omega \rightarrow \text{Bool}) \rightarrow (\Sigma^\omega \rightarrow \text{Bool})$ is defined by*

$$\{p\}(q) = p \wedge q.$$

The assert transformer $\{p\}$ can be seen as modeling a system which accepts all input traces that satisfy p , and rejects all others. For all the traces that it accepts, the system simply transfers them, i.e., it behaves as the identity function.

To express MPTs such as assert transformers syntactically, let us introduce some notation. First, we can use lambda notation for predicates, as in $\lambda(\sigma, \sigma') : (\sigma = \sigma')$ for some predicate $p : \Sigma^\omega \rightarrow \Sigma^\omega \rightarrow \mathbb{B}$ which returns true whenever it receives two equal traces. Then, instead of writing $\{\lambda(\sigma, \sigma') : (\sigma = \sigma')\}$ for the corresponding assert transformer $\{p\}$, we will use the slightly lighter notation $\{\sigma, \sigma' \mid \sigma = \sigma'\}$.

Definition 11 (Demonic update). *Let $r : \Sigma_x^\omega \rightarrow \Sigma_y^\omega \rightarrow \text{Bool}$ be a relation. The demonic update property transformer $[r] : (\Sigma_y^\omega \rightarrow \text{Bool}) \rightarrow (\Sigma_x^\omega \rightarrow \text{Bool})$ is defined by*

$$[r](q) = \{\sigma \mid \forall \sigma' : r(\sigma)(\sigma') \rightarrow \sigma' \in q\}.$$

That is, $[r](q)$ contains all input traces σ which are guaranteed to result into an output trace in q when fed into the (generally non-deterministic) input-output relation r . The term “demonic update” comes from the refinement calculus literature [10].

Similarly to assert, we introduce a lightweight notation for the demonic update. If r is an expression in σ and σ' , then $[\sigma \rightsquigarrow \sigma' \mid r] = [\lambda(\sigma, \sigma') : r]$. For example, $[\sigma_x, \sigma_y \rightsquigarrow \sigma_z \mid \forall i : \sigma_z(i) = \sigma_x(i) + \sigma_y(i)]$ is the system which produces as output the sequence $\sigma_z = (\lambda i : \sigma_x(i) + \sigma_y(i))$, where σ_x and σ_y are the input sequences. If e is an expression in σ , then $[\sigma \rightsquigarrow e]$ is defined to be $[\sigma \rightsquigarrow \sigma' \mid \sigma' = e]$, where σ' is a new variable different from σ and which does not occur free in e . For example, $[\sigma \rightsquigarrow (\lambda i : \sigma(i) + 1)] = [\sigma \rightsquigarrow \sigma' \mid \sigma' = (\lambda i : \sigma(i) + 1)]$.

The following lemma states that Skip can be defined as an assert transformer, or as a demonic update transformer.

Lemma 2. $\text{Skip} = [\sigma \rightsquigarrow \sigma' \mid \sigma = \sigma'] = \{\top\} = \{\sigma \mid \text{true}\}$.

In general `Skip`, `Fail`, and other property transformers are polymorphic with respect to their input and output types. In `Skip` the input and output types must be the same. `Fail`, on the other hand, may have an input type and a different output type.

Definition 12 (Angelic update). *Let $r : \Sigma_x^\omega \rightarrow \Sigma_y^\omega \rightarrow \text{Bool}$ be a relation. The angelic update property transformer $\{r\} : (\Sigma_y^\omega \rightarrow \text{Bool}) \rightarrow (\Sigma_x^\omega \rightarrow \text{Bool})$ is defined by*

$$\{r\}(q) = \{\sigma \mid \exists \sigma' : r(\sigma)(\sigma') \wedge \sigma' \in q\}.$$

An input sequence σ is in $\{r\}(q)$ if there exists an output sequence σ' such that $r(\sigma)(\sigma')$ and $\sigma' \in q$. Notice the duality between the angelic and demonic update transformers. Consider, for example, a relation $r = \{(\sigma, \sigma'), (\sigma, \sigma'')\}$. If $q = \{\sigma', \sigma''\}$, then $\{r\}(q) = [r](q) = \{\sigma\}$. If $q = \{\sigma'\}$ then $\{r\}(q) = \{\sigma\}$, while $[r](q) = \emptyset$.

We use a lightweight notation for the angelic update transformer, similar to the one for demonic update. If r is an expression in σ and σ' , then $\{\sigma \rightsquigarrow \sigma' \mid r\} = \{\lambda(\sigma, \sigma') : r\}$.

Lemma 3. *Assert is a particular case of angelic update: $\{p\} = \{\sigma \rightsquigarrow \sigma' \mid p(\sigma) \wedge \sigma = \sigma'\}$.*

4.1.2 Operators on MPTs: Function Composition, Product, and Fusion

As we shall see in §4.3, the semantics of composition operators in the language of components will be defined by corresponding composition operators on MPTs. We now introduce the latter operators on MPTs. First, we begin by the operators that have been known in the literature, and are recalled here. In §4.1.3 we introduce some novel operators explicitly designed in order to handle feedback composition.

Serial composition of MPTs (and property transformers in general) is simply function composition. Let $S : (\Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$ and $T : (\Sigma_z^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_y^\omega \rightarrow \mathbb{B})$ be two property transformers. Then $S \circ T : (\Sigma_z^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$, is the function composition of S and T , i.e., $\forall q : (S \circ T)(q) = S(T(q))$. Note that serial composition preserves monotonicity, so that if S and T are MPTs, then $S \circ T$ is also an MPT. Also note that `Skip` is the neutral element for serial composition, i.e., $S \circ \text{Skip} = \text{Skip} \circ S = S$.

To express parallel composition of components, we need a kind of Cartesian product operation on property transformers. We define such an operation below. Similar operations for predicate transformers have been proposed in [9].

Definition 13 (Product). *Let $S : (\Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$ and $T : (\Sigma_v^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_u^\omega \rightarrow \mathbb{B})$. The product of S and T , denoted $S \otimes T : (\Sigma_y^\omega \times \Sigma_v^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \times \Sigma_u^\omega \rightarrow \mathbb{B})$, is given by*

$$(S \otimes T)(q) = \{(\sigma, \sigma') \mid \exists p : \Sigma_y^\omega \rightarrow \mathbb{B}, p' : \Sigma_v^\omega \rightarrow \mathbb{B} : p \times p' \subseteq q \wedge \sigma \in S(p) \wedge \sigma' \in T(p')\}$$

where $(p \times p')(\sigma_y, \sigma_v) = p(\sigma_y) \wedge p'(\sigma_v)$.

Lemma 4. *For arbitrary S and T , $S \otimes T$ is monotonic.*

The neutral element for the product composition is the `Skip` MPT that has `Unit` as input and output type.

In order to define a feedback operation on MPTs, we first define two auxiliary operations: *Fusion* and *IterateOmega*. Fusion is an extension of a similar operator introduced previously for predicate transformers in [9]. IterateOmega is a novel operator introduced in the sequel.

Definition 14 (Fusion). *If $S = \{S_i\}_{i \in I}$, $S_i : (\Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$ is a collection of MPTs, then the fusion of S is the MPT $\text{Fusion}_{i \in I}(S_i) : (\Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$ defined by*

$$(\text{Fusion}_{i \in I}(S_i))(q) = \{\sigma \mid \exists p : I \rightarrow \Sigma_y^\omega \rightarrow \mathbb{B} : \bigcap_i p_i \subseteq q \wedge \sigma \in \bigcap_i S_i(p_i)\}$$

The `Fusion` operator satisfies the following property.

Lemma 5. *For $I \neq \emptyset$ we have*

$$\text{Fusion}_{i \in I}(\{p_i\} \circ [r_i]) = \{\bigcap_{i \in I} p_i\} \circ [\bigcap_{i \in I} r_i].$$

4.1.3 Novel Operators Used in Semantical Definition of Feedback

The `IterateOmega` operator is defined as follows:

Definition 15 (`IterateOmega`).

$$\text{IterateOmega}(S) = \text{Fusion}_{n \in \mathbb{N}}(S^n \circ [\sigma \rightsquigarrow \sigma' \mid \forall i : i + 1 < n \rightarrow \sigma_i = \sigma'_i])$$

The *feedback* operator consists of connecting the first output of an MPT S with its first input. Formally, feedback is defined as follows.

Definition 16 (Feedback composition of an MPT). *Let $S : (\Sigma_u \times \Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_u \times \Sigma_x^\omega \rightarrow \mathbb{B})$ be an MPT. The feedback operator on S , denoted `Feedback`(S), is given by the MPT*

$$\begin{aligned} \text{Feedback}(S) = & \{\sigma_x \rightsquigarrow \sigma_x, \sigma_u, \sigma_y\} \\ & \circ \text{IterateOmega}([\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_x, \sigma_u, \sigma_x] \circ (\text{Skip} \otimes S)) \\ & \circ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_y] \end{aligned}$$

Example. As an example we show how to derive `Feedback`(S) for $S = [\sigma_u, \sigma_x \rightsquigarrow 0 \cdot \sigma_x + 0 \cdot \sigma_u, 0 \cdot \sigma_x + 0 \cdot \sigma_u]$, where $\sigma + \sigma' = (\lambda i : \sigma(i) + \sigma'(i))$, and $0 \cdot \sigma$ is 0 concatenated with σ . For now, we note that S is the semantics of the composite component `Add`; `UnitDelay`; `Split`, which corresponds to the inner part of the diagram of Fig. 2, before applying feedback. We will complete the formal definition of the the semantics of this diagram in §4.3. For now, we focus on deriving `Feedback`(S), in order to illustrate how the `Feedback` operator works.

Let

$$T = [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_x, \sigma_u, \sigma_x] \circ (\text{Skip} \otimes S) = [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_x, 0 \cdot \sigma_x + 0 \cdot \sigma_u, 0 \cdot \sigma_x + 0 \cdot \sigma_u]$$

Then, we have

$$\begin{aligned} T \circ T = & [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_x, 0 \cdot \sigma_x + 0 \cdot 0 \cdot \sigma_x + 0 \cdot 0 \cdot \sigma_u, 0 \cdot \sigma_x + 0 \cdot 0 \cdot \sigma_x + 0 \cdot 0 \cdot \sigma_u] \\ \dots \\ T^n = & [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_x, 0 \cdot \sigma_x + \dots + 0^n \cdot \sigma_x + 0^n \cdot \sigma_u, 0 \cdot \sigma_x + \dots + 0^n \cdot \sigma_x + 0^n \cdot \sigma_u] \end{aligned}$$

where 0^n is a finite sequence of n 0s. We also have

$$\begin{aligned} T^n \circ [\sigma \rightsquigarrow \sigma' \mid \forall i : i + 1 < n \rightarrow \sigma_i = \sigma'_i] \\ = \\ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma \mid \forall i : i + 1 < n \rightarrow \sigma(i) = (\sigma_x(i), \Sigma_{j < i} \sigma_x(j), \Sigma_{j < i} \sigma_x(j))] \end{aligned}$$

Then

$$\begin{aligned} \text{Feedback}(S) &= \\ &= \{\sigma_x \rightsquigarrow \sigma_x, \sigma_u, \sigma_y\} \circ \text{IterateOmega}(T) \circ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_y] \\ &= \{\sigma_x \rightsquigarrow \sigma_x, \sigma_u, \sigma_y\} \circ \text{Fusion}_{n \in \mathbb{N}}(T^n \circ [\sigma \rightsquigarrow \sigma' \mid \forall i : i + 1 < n \rightarrow \sigma_i = \sigma'_i]) \circ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_y] \\ &= \{\sigma_x \rightsquigarrow \sigma_x, \sigma_u, \sigma_y\} \circ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma \mid \forall i : \sigma(i) = (\sigma_x(i), \Sigma_{j < i} \sigma_x(j), \Sigma_{j < i} \sigma_x(j))] \circ [\sigma_x, \sigma_u, \sigma_y \rightsquigarrow \sigma_y] \\ &= [\sigma_x \rightsquigarrow \sigma_y \mid \forall i : \sigma_y(i) = \Sigma_{j < i} \sigma_x(j)] \end{aligned}$$

Finally we obtain

$$\text{Feedback}(S) = [\sigma_x \rightsquigarrow \sigma_y \mid \forall i : \sigma_y(i) = \Sigma_{j < i} \sigma_x(j)] \quad (1)$$

This is the system that outputs the trace $(\lambda i : \Sigma_{j < i} \sigma_x(j))$ for input trace σ_x .

4.1.4 Refinement

A key element of RCRS, as of other compositional frameworks, is the notion of refinement, which enables substitutability and other important concepts of compositionality. Semantically, refinement is defined as follows:

Definition 17 (Refinement). *Let $S, T : (\Sigma_y^\omega \rightarrow \mathbb{B}) \rightarrow (\Sigma_x^\omega \rightarrow \mathbb{B})$ be two MPTs. We say that T refines S (or that S is refined by T), written $S \sqsubseteq T$, if and only if $\forall q : S(q) \subseteq T(q)$.*

All operations introduced on MPTs preserve the refinement relation:

Theorem 1. *If S, T, S', T' are MPTs of appropriate types such that $S \sqsubseteq S'$ and $T \sqsubseteq T'$, then*

1. $S \circ T \sqsubseteq S' \circ T'$ and $S \otimes T \sqsubseteq S' \otimes T'$ and $\text{Fusion}(S, T) \sqsubseteq \text{Fusion}(S', T')$, [9, 10]
2. $\text{IterateOmega}(S) \sqsubseteq \text{IterateOmega}(S')$
3. $\text{Feedback}(S) \sqsubseteq \text{Feedback}(S')$

4.2 Subclasses of MPTs

Monotonic property transformers are a very rich and powerful class of semantic objects. In practice, the systems that we deal with often fall into restricted subclasses of MPTs, which are easier to represent syntactically and manipulate symbolically. We introduce these subclasses next.

4.2.1 Relational MPTs

Definition 18 (Relational property transformers). *A relational property transformer S is an MPT of the form $\{p\} \circ [r]$. We call p the precondition of S and r the input-output relation of S .*

Relational property transformers correspond to *conjunctive* transformers [10]. A transformer S is conjunctive if it satisfies $S(\bigcap_{i \in I} q_i) = \bigcap_{i \in I} S(q_i)$ for all $(q_i)_{i \in I}$ and $I \neq \emptyset$.

`Fail`, `Skip`, any assert transformer $\{p\}$, and any demonic update transformer $[r]$, are RPTs. Indeed, `Fail` can be written as $\{\sigma \mid \text{false}\} \circ [\sigma \rightsquigarrow \sigma' \mid \text{true}]$. `Skip` can be written as $\{\sigma \mid \text{true}\} \circ [\sigma \rightsquigarrow \sigma]$. The assert transformer $\{p\}$ can be written as the RPT $\{p\} \circ [\sigma \rightsquigarrow \sigma]$. Finally, the demonic update transformer $[r]$ can be written as the RPT $\{\sigma \mid \text{true}\} \circ [r]$. Angelic update transformers are generally not RPTs: the angelic update transformer $\{\sigma \rightsquigarrow \sigma' \mid \text{true}\}$ is not an RPT, as it is not conjunctive.

Examples. Suppose we wish to specify a system that performs division. Here are two possible ways to represent this system with RPTs:

$$S_1 = \{\top\} \circ [\sigma_x, \sigma_y \rightsquigarrow \sigma_z \mid \forall i : \sigma_y(i) \neq 0 \wedge \sigma_z(i) = \frac{\sigma_x(i)}{\sigma_y(i)}]$$

$$S_2 = \{\sigma_x, \sigma_y \mid \forall i : \sigma_y(i) \neq 0\} \circ [\sigma_x, \sigma_y \rightsquigarrow \sigma_z \mid \sigma_z(i) = \frac{\sigma_x(i)}{\sigma_y(i)}]$$

Although S_1 and S_2 are both relational, they are not equivalent transformers. S_1 is input-receptive: it accepts all input traces. However, if at some step i the input $\sigma_y(i)$ is 0, then the output $\sigma_z(i)$ is arbitrary (non-deterministic). In contrast, S_2 is non-input-receptive as it accepts only those traces σ_y that are guaranteed to be non-zero at every step, i.e., those that satisfy the condition $\forall i : \sigma_y(i) \neq 0$.

Theorem 2 (RPTs are closed under serial, product and fusion compositions). *Let $S = \{p\} \circ [r]$ and $S' = \{p'\} \circ [r']$ be two RPTs, with p, p', r and r' of appropriate types. Then*

$$S \circ S' = \{\sigma \mid p(\sigma) \wedge (\forall \sigma' : r(\sigma)(\sigma') \rightarrow p'(\sigma'))\} \circ [r \circ r']$$

and

$$S \otimes S' = \{\sigma_x, \sigma_y \mid p(\sigma_x) \wedge p'(\sigma_y)\} \circ [\sigma_x, \sigma_y \rightsquigarrow \sigma'_x, \sigma'_y \mid r(\sigma_x)(\sigma'_x) \wedge r'(\sigma_y)(\sigma'_y)]$$

and

$$\text{Fusion}(S, S') = \{p \wedge p'\} \circ [r \wedge r']$$

RPTs are not closed under the feedback operation. Indeed, we have

$$\text{Feedback}([\sigma_u, \sigma_x \rightsquigarrow \sigma_u, \sigma_u]) = \{\sigma_x \rightsquigarrow \sigma_y \mid \text{true}\}$$

which is a non-relational angelic update transformer as we said above.

Next theorem shows that the refinement of RPTs can be reduced to proving a first order property.

Theorem 3. *For p, p', r, r' of appropriate types we have.*

$$\{p\} \circ [r] \sqsubseteq \{p'\} \circ [r'] \iff (p \subseteq p' \wedge (\forall \sigma_x : p(\sigma_x) \Rightarrow r(\sigma_x) \subseteq r'(\sigma_x)))$$

4.2.2 Guarded MPTs

Relational property transformers correspond to systems that have natural syntactic representations, as the composition $\{p\} \circ [r]$ where the predicate p and the relation r can be represented syntactically in some logic. Unfortunately, RPTs are still too powerful. In particular, they allow system semantics that cannot be implemented. For example, consider the RPT $\text{Magic} = [\sigma \rightsquigarrow \sigma' \mid \text{false}]$. It can be shown that for any output property q (including \perp), we have $\text{Magic}(q) = \top$. Recall that, viewed as a set, \top is the set of all traces. This means that, no matter what the post-condition q is, Magic somehow manages to produce output traces satisfying q no matter what the input trace is (hence the name “magic”). In general, an MPT S is said to be *non-miraculous* (or to satisfy the *law of excluded miracle*) if $S(\perp) = \perp$. We note that in [28], sequential programs are modeled using predicate transformers that are *conjunctive* and satisfy the *law of excluded miracle*.

We want to further restrict RPTs so that miraculous behaviour does not arise. Specifically, for an RPT $S = \{p\} \circ [r]$ and an input sequence σ , if there is no σ' such that $r(\sigma)(\sigma')$ is satisfied, then we want σ to be illegal, i.e., we want $p(\sigma) = \text{false}$. We can achieve this by taking p to be $\text{in}(r)$. Recall that if $r : X \rightarrow Y \rightarrow \mathbb{B}$, then $\text{in}(r)(x) = (\exists y : r(x)(y))$. Taking p to be $\text{in}(r)$ effectively means that p and r are combined into a single specification r which can also restrict the inputs. This is also the approach followed in the theory of relational interfaces [77].

Definition 19 (Guarded property transformers). *The guarded property transformer (GPT) of a relation r is an RPT, denoted $\{r\} = \{\text{in}(r)\} \circ [r]$.*

It can be shown that an MPT S is a GPT if and only if S is conjunctive and non-miraculous [10]. Fail , Skip , and any assert property transformer are GPTs. Indeed, $\text{Fail} = \{\perp\}$ and $\text{Skip} = \{\sigma \rightsquigarrow \sigma \mid \top\}$. The assert transformer can be written as $\{p\} = \{\sigma \rightsquigarrow \sigma \mid p(\sigma)\}$. The angelic and demonic update property transformers are generally not GPTs. The angelic update property transformer is not always conjunctive in order to be a GPT. The demonic update property transformer is not in general a GPT because is not always non-miraculous ($\text{Magic}(\perp) = \top \neq \perp$). The demonic update transformer $[r]$ is a GPT if and only if $\text{in}(r) = \text{true}$ and in this case we have $[r] = \{r\}$.

Theorem 4 (GPTs are closed under serial and parallel compositions). *Let $S = \{r\}$ and $S' = \{r'\}$ be two GPTs with r and r' of appropriate types. Then*

$$S \circ S' = \{\sigma_x \rightsquigarrow \sigma_z \mid \text{in}(r)(\sigma_x) \wedge (\forall \sigma_y : r(\sigma_x)(\sigma_y) \rightarrow \text{in}(r')(\sigma_y)) \wedge (r \circ r')(\sigma_x, \sigma_z)\}$$

and

$$S \otimes S' = \{\sigma_x, \sigma_y \rightsquigarrow \sigma'_x, \sigma'_y \mid r(\sigma_x)(\sigma'_x) \wedge r'(\sigma_y)(\sigma'_y)\}$$

GPTs are not closed under Fusion and Feedback operations. Indeed, we have seen in the previous section that the feedback operator applied to the GPT $[\sigma_u, \sigma_x \rightsquigarrow \sigma_u, \sigma_u]$ is not an RPT, and therefore it is not a GPT either. For the fusion operator, we have $\text{Fusion}([x \rightsquigarrow 0], [x \rightsquigarrow 1]) = [\perp]$, which is not a GPT.

4.2.3 Other subclasses and overview

The containment relationships among the various subclasses of MPTs are illustrated in Fig. 4. In addition to the subclasses discussed above, we introduce several more subclasses of MPTs in the sections that follow, when we assign semantics (in terms of MPTs) to the various atomic components in our component language. For instance, QLTL components give rise to *QLTL property transformers*. Similarly, STS components, stateless STS components, etc., give rise to corresponding subclasses of MPTs. The containment relationships between these classes will be proven in the sections that follow. For ease of reference, we provide some forward links to these results also here. The fact that QLTL property transformers are GPTs follows by definition of the semantics of QLTL components: see §4.3, equation (2). The fact that STS property transformers are a special case of QLTL property transformers follows from the transformation of an STS component into a semantically equivalent QLTL component: see §5.1 and Theorem 7. The inclusions for subclasses of STS property transformers follow by definition of the corresponding components (see also Fig. 5).

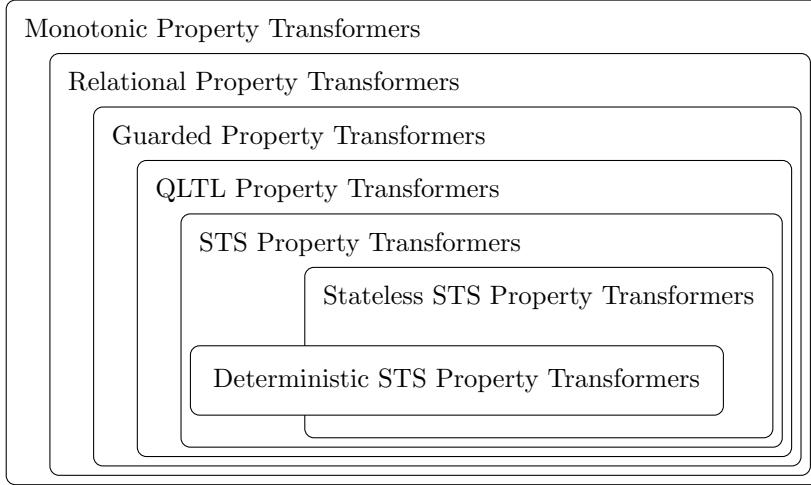


Figure 4: Overview of the property transformer classes and their containment relations.

4.3 Semantics of Components as MPTs

We are now ready to define the semantics of our language of components in terms of MPTs. Let C be a well formed component. The semantics of C , denoted $\llbracket C \rrbracket$, is a property transformer of the form:

$$\llbracket C \rrbracket : ((\Sigma_{out}(C))^\omega \rightarrow \mathbb{B}) \rightarrow ((\Sigma_{in}(C))^\omega \rightarrow \mathbb{B}).$$

We define $\llbracket C \rrbracket$ by induction on the structure of C . First we give the semantics of QLTL components and composite components:

$$\llbracket \text{qltl}(x, y, \varphi) \rrbracket = \{\sigma_x \rightsquigarrow \sigma_y \mid (\sigma_x, \sigma_y) \models \varphi\} \quad (2)$$

$$\llbracket C ; C' \rrbracket = \llbracket C \rrbracket \circ \llbracket C' \rrbracket \quad (3)$$

$$\llbracket C \parallel C' \rrbracket = \llbracket C \rrbracket \otimes \llbracket C' \rrbracket \quad (4)$$

$$\llbracket \text{fdbk}(C) \rrbracket = \text{Feedback}(\llbracket C \rrbracket) \quad (5)$$

The semantics of QLTL components satisfies the following property:

Lemma 6. *If φ is a LTL formula on variables x and y , we have:*

$$\llbracket \text{qltl}(x, y, (\exists y : \varphi) \wedge \varphi) \rrbracket = \llbracket \text{qltl}(x, y, \varphi) \rrbracket$$

To define the semantics of STS components, we first introduce some auxiliary notation. Consider an STS component $C = \text{sts}(x, y, s, \text{init_exp}, \text{trs_exp})$. We define the predicate $\text{run}_C : \Sigma_s^\omega \times \Sigma_x^\omega \times \Sigma_y^\omega \rightarrow \mathbb{B}$ as

$$\text{run}_C(\sigma_s, \sigma_x, \sigma_y) = (\forall i : \text{trs_exp}(\sigma_s(i), \sigma_x(i))(\sigma_s(i+1), \sigma_y(i))).$$

Intuitively, if $\sigma_x \in \Sigma_x^\omega$ is the input sequence, $\sigma_y \in \Sigma_y^\omega$ is the output sequence, and $\sigma_s \in \Sigma_s^\omega$ is the sequence of values of state variables, then $\text{run}_C(\sigma_s, \sigma_x, \sigma_y)$ holds if at each step of the execution, the current state, current input, next state, and current output, satisfy the trs_exp predicate.

We also formalize the illegal input traces of STS component C as follows:

$$\begin{aligned} \text{illegal}_C(\sigma_x) = & (\exists \sigma_s, \sigma_y, k : \text{init_exp}(\sigma_s(0)) \wedge (\forall i < k : \text{trs_exp}(\sigma_s(i), \sigma_x(i))(\sigma_s(i+1), \sigma_y(i))) \wedge \\ & \neg \text{in}(\text{trs_exp})(\sigma_s(k), \sigma_x(k))) \end{aligned}$$

Essentially, $\text{illegal}_C(\sigma_x)$ states that there exists some point in the execution where the current state and current input violate the precondition $\text{in}(\text{trs_exp})$ of predicate trs_exp , i.e., there exist no output and next state to satisfy trs_exp for that given current state and input.

Then, the semantics of STS component C is given by:

$$\llbracket C \rrbracket = \{\neg \text{illegal}_C\} \circ [\sigma_x \rightsquigarrow \sigma_y \mid (\exists \sigma_s : \text{init_exp}(\sigma_s(0)) \wedge \text{run}_C(\sigma_s, \sigma_x, \sigma_y))] \quad (6)$$

We give semantics to stateless and/or deterministic STS components using the corresponding mappings from general STS components. If C is a stateless STS, C' is a deterministic STS, and C'' is a stateless deterministic STS, then:

$$\llbracket C \rrbracket = \llbracket \text{stateless2sts}(C) \rrbracket \quad (7)$$

$$\llbracket C' \rrbracket = \llbracket \text{det2sts}(C') \rrbracket \quad (8)$$

$$\llbracket C'' \rrbracket = \llbracket \text{stateless_det2det}(C'') \rrbracket \quad (9)$$

4.3.1 Example: Two Alternative Derivations of the Semantics of Diagram Sum of Fig. 2

To illustrate our semantics, we provide two alternative derivations of the semantics of the `Sum` system of Fig. 2.

First, let us consider `Sum` as a composite component:

$$\text{Sum} = \text{fdbk}(\text{Add}; \text{UnitDelay}; \text{Split})$$

where

$$\begin{aligned} \text{Add} &= \text{stateless_det}((u, x), \text{true}, u + x) \\ \text{UnitDelay} &= \text{det}(x, s, 0, \text{true}, x, s) \\ \text{Split} &= \text{stateless_det}(x, \text{true}, (x, x)) \end{aligned}$$

We have

$$\begin{aligned} \llbracket \text{Sum} \rrbracket &= \llbracket \text{fdbk}(\text{Add}; \text{UnitDelay}; \text{Split}) \rrbracket \\ &= \text{Feedback}(\llbracket \text{Add} \rrbracket \circ \llbracket \text{UnitDelay} \rrbracket \circ \llbracket \text{Split} \rrbracket) \end{aligned}$$

For `Add`, `UnitDelay` and `Split`, all inputs are legal, so $\text{illegal}_C = \perp$ for all $C \in \{\text{Add}, \text{UnitDelay}, \text{Split}\}$. After simplifications, we get:

$$\begin{aligned} \llbracket \text{Add} \rrbracket &= [\sigma_u, \sigma_x \rightsquigarrow \sigma_u + \sigma_x] \\ \llbracket \text{UnitDelay} \rrbracket &= [\sigma_x \rightsquigarrow 0 \cdot \sigma_x] \\ \llbracket \text{Split} \rrbracket &= [\sigma_x \rightsquigarrow \sigma_x, \sigma_x] \end{aligned}$$

The semantics of **Sum** is given by

$$\begin{aligned}
& \llbracket \text{Sum} \rrbracket \\
&= \text{Feedback}(\llbracket \text{Add} \rrbracket \circ \llbracket \text{UnitDelay} \rrbracket \circ \llbracket \text{Split} \rrbracket) \\
&= \text{Feedback}([\sigma_u, \sigma_x \rightsquigarrow \sigma_u + \sigma_x] \circ [\sigma_x \rightsquigarrow 0 \cdot \sigma_x] \circ [\sigma_x \rightsquigarrow \sigma_x, \sigma_x]) \\
&= \text{Feedback}([\sigma_u, \sigma_x \rightsquigarrow 0 \cdot \sigma_x + 0 \cdot \sigma_u, 0 \cdot \sigma_x + 0 \cdot \sigma_u]) \\
&= \{\text{Using (1)}\} \\
&\quad [\sigma_x \rightsquigarrow \sigma_y \mid \forall i : \sigma_y(i) = \sum_{j < i} \sigma_x(j)]
\end{aligned}$$

We obtain:

$$\llbracket \text{Sum} \rrbracket = [\sigma_x \rightsquigarrow \sigma_y \mid \forall i : \sigma_y(i) = \sum_{j < i} \sigma_x(j)]. \quad (10)$$

Next, let us assume that the system has been characterized already as an atomic component:

$$\text{SumAtomic} = \text{sts}(x, y, s, s = 0, y = s \wedge s' = s + x).$$

The semantics of **SumAtomic** is given by

$$\llbracket \text{SumAtomic} \rrbracket = \{\neg \text{illegal}_{\text{SumAtomic}}\} \circ [\sigma_x \rightsquigarrow \sigma_y \mid \exists \sigma_s : \sigma_s(0) = 0 \wedge \text{run}_{\text{SumAtomic}}(\sigma_s, \sigma_x, \sigma_y)]$$

where $\text{illegal}_{\text{SumAtomic}} = \perp$ because there are no restrictions on the inputs of **SumAtomic**, and

$$\text{run}_{\text{SumAtomic}}(\sigma_s, \sigma_x, \sigma_y) = (\forall i : \sigma_y(i) = \sigma_s(i) \wedge \sigma_s(i+1) = \sigma_s(i) + \sigma_x(i))$$

We have

$$\llbracket \text{SumAtomic} \rrbracket = [\sigma_x \rightsquigarrow \sigma_y \mid \exists \sigma_s : \sigma_s(0) = 0 \wedge (\forall i : \sigma_y(i) = \sigma_s(i) \wedge \sigma_s(i+1) = \sigma_s(i) + \sigma_x(i))]$$

which is equivalent to (10).

4.3.2 Characterization of Legal Input Traces

The following lemma characterizes legal input traces for various types of MPTs:

Lemma 7. *The set of legal input traces of an RPT $\{p\} \circ [r]$ is p :*

$$\text{legal}(\{p\} \circ [r]) = p.$$

The set of legal input traces of a GPT $[r]$ is $\text{in}(r)$:

$$\text{legal}([r]) = \text{in}(r).$$

The set of legal input traces of an STS component $C = \text{sts}(x, y, s, \text{init}, r)$ is equal to $\neg \text{illegal}_C$:

$$\text{legal}(\llbracket C \rrbracket) = \neg \text{illegal}_C.$$

The set of legal input traces of a QLTL component $C = \text{qltl}(x, y, \varphi)$ is:

$$\text{legal}(\llbracket C \rrbracket) = \{\sigma_x \mid \sigma_x \models \exists y : \varphi\}.$$

4.3.3 Semantic Equivalence and Refinement for Components

Definition 20. *Two components C and C' are (semantically) equivalent, denoted $C \equiv C'$, if $\llbracket C \rrbracket = \llbracket C' \rrbracket$. Component C is refined by component C' , denoted $C \preceq C'$, if $\llbracket C \rrbracket \sqsubseteq \llbracket C' \rrbracket$.*

The relation \equiv is an equivalence relation, and \preceq is a preorder relation (i.e., reflexive and transitive). We also have

$$(C \preceq C' \wedge C' \preceq C) \iff (C \equiv C')$$

4.3.4 Compositionality Properties

Several desirable compositionality properties follow from our semantics:

Theorem 5. *Let C_1, C_2, C_3 , and C_4 be four (possibly composite) components. Then:*

1. (Serial composition is associative:) $(C_1 ; C_2) ; C_3 \equiv C_1 ; (C_2 ; C_3)$.
2. (Parallel composition is associative:) $(C_1 \parallel C_2) \parallel C_3 \equiv C_1 \parallel (C_2 \parallel C_3)$.
3. (Parallel composition distributes over serial composition:) $(C_1 \parallel C_2) ; (C_3 \parallel C_4) \equiv (C_1 ; C_3) \parallel (C_2 ; C_4)$.
4. (Refinement is preserved by composition:) If $C_1 \preceq C_2$ and $C_3 \preceq C_4$, then:
 - (a) $C_1 ; C_3 \preceq C_2 ; C_4$
 - (b) $C_1 \parallel C_3 \preceq C_2 \parallel C_4$
 - (c) $\mathbf{fdbk}(C_1) \preceq \mathbf{fdbk}(C_2)$

In addition to the above, requirements that a component satisfies are preserved by refinement. Informally, if C satisfies some requirement φ and $C \preceq C'$ then C' also satisfies φ . Although we have not formally defined what requirements are and what it means for a component to satisfy a requirement, these concepts are naturally captured in the RCRS framework via the semantics of components as MPTs. In particular, since our components are generally *open* systems (i.e., they have inputs), we can express requirements using *Hoare triples* of the form $p\{C\}q$, where C is a component, p is an input property, and q is an output property. Then, $p\{C\}q$ holds iff the outputs of C are guaranteed to satisfy q provided the inputs of C satisfy p . Formally: $p\{C\}q \iff p \subseteq \llbracket C \rrbracket(q)$.

Theorem 6. $C \preceq C'$ iff $\forall p, q : p\{C\}q \rightarrow p\{C'\}q$.

Theorem 6 shows that refinement is equivalent to *substitutability*. Substitutability states that a component C' can replace another component C in any context, i.e., $\forall p, q : p\{C\}q \rightarrow p\{C'\}q$.

5 Symbolic Reasoning

So far we have defined the syntax and semantics of RCRS. These already allow us to specify and reason about systems in a compositional manner. However, such reasoning is difficult to do “by hand”. For example, if we want to check whether a component C is refined by another component C' , we must resort to proving the refinement relation $\llbracket C \rrbracket \sqsubseteq \llbracket C' \rrbracket$ of their corresponding MPTs, $\llbracket C \rrbracket$ and $\llbracket C' \rrbracket$. This is not an easy task, as MPTs are complex mathematical objects. Instead, we would like to have computer-aided, and ideally fully automatic techniques. In the above example of checking refinement, for instance, we would like ideally to have an algorithm that takes as input the syntactic descriptions of C and C' and replies yes/no based on whether $\llbracket C \rrbracket \sqsubseteq \llbracket C' \rrbracket$ holds. We say “ideally” because we know that in general such an algorithm cannot exist. This is because we are not making a-priori any restrictions on the logics used to describe C and C' , which means that the existence of an algorithm will depend on decidability of these logics. In this section, we describe how reasoning in RCRS can be done *symbolically*, by automatically manipulating the formulas used to specify the components involved. As we shall show, most problems can be reduced to checking satisfiability of first-order formulas formed by combinations of the formulas of the original components. This means that the problems are decidable whenever the corresponding first-order logics are decidable.

5.1 Symbolic Transformation of STS Components to QLTL Components

Our framework allows the specification of several types of atomic components, some of which are special cases of others, as summarized in Fig. 5. In §3, we have already shown how the different types of STS components are related, from the most specialized deterministic stateless STS components, to the general

STS components. By definition, the semantics of the special types of STS components are defined via the semantics of general STS components (see §4). In this subsection, we show that STS components can be viewed as special cases of QLTL components.

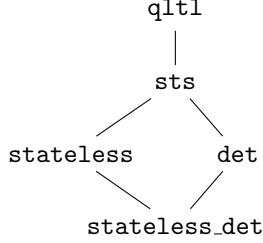


Figure 5: Lattice of atomic components: lower types are special cases of higher types.

Specifically, we show how an STS component can be transformed into a semantically equivalent QLTL component. This transformation also shows that STS property transformers are a special case of QLTL property transformers, as already claimed in Fig. 4. Note that this containment is not obvious simply by looking at the definitions of these MPT subclasses (c.f. §4.3), as QLTL property transformers are defined as GPTs (equation 2), whereas STS property transformers are defined as RPTs (equation 6). Although RPTs are generally a superclass, not a subclass, of GPTs, the transformation proposed below shows that the RPTs obtained by STS components can indeed be captured as GPTs. The transformation of STS into QLTL components also enables us to apply several algorithms which are available for QLTL components, to STS components as well.

Consider an STS component $C = \text{sts}(x, y, s, \text{init}, \text{trs})$. We can transform C into a QLTL component using the mapping sts2qltl :

$$\text{sts2qltl}(C) = \text{qltl}(x, y, (\forall s, y : \text{init} \rightarrow (\varphi \mathbf{L} \varphi')) \wedge (\exists s : \text{init} \wedge \mathbf{G} \varphi)) \quad (11)$$

where $\varphi = \text{trs}[s' := \circlearrowleft s]$ and $\varphi' = (\exists s', y : \text{trs})$.

The theorem that follows demonstrates the correctness of the above transformation:

Theorem 7. *For any STS component C , $C \equiv \text{sts2qltl}(C)$.*

Example. It is instructive to see how the above transformation specializes to some special types of STS components. In particular, we will show how it specializes to stateless STS components.

Let $C = \text{stateless}(x, y, \text{trs})$ and let $C' = \text{stateless2sts}(C) = \text{sts}(x, y, (), \text{true}, \text{trs})$. Applying the sts2qltl transformation to C' , for which $s = ()$ and $\text{init} = \text{true}$, we obtain:

$$\text{sts2qltl}(C') = \text{qltl}(x, y, (\forall y : (\varphi \mathbf{L} \varphi') \wedge \mathbf{G} \varphi))$$

where $\varphi = \text{trs}[() := \circlearrowleft ()] = \text{trs}$ and $\varphi' = (\exists y : \text{trs})$. Using the properties of Lemma 1, and the fact that semantically equivalent LTL formulas result in semantically equivalent QLTL components, we can simplify

`sts2qltl`(C') further:

$$\begin{aligned}
& \text{sts2qltl}(C') \\
= & \{\text{Definition of } \text{sts2qltl}\} \\
& \text{qltl}(x, y, (\forall y : (\text{trs } \mathbf{L} (\exists y : \text{trs}))) \wedge \mathbf{G} \text{ trs}) \\
\equiv & \{\text{Lemma 1, } \text{trs} \text{ does not contain temporal operators, and } y \text{ is not free in } (\exists y : \text{trs})\} \\
& \text{qltl}(x, y, ((\exists y : \text{trs}) \mathbf{L} (\exists y : \text{trs})) \wedge \mathbf{G} \text{ trs}) \\
\equiv & \{\text{Lemma 1}\} \\
& \text{qltl}(x, y, \mathbf{G} (\exists y : \text{trs}) \wedge \mathbf{G} \text{ trs}) \\
\equiv & \{\text{Lemma 1 and } \text{trs} \text{ does not contain temporal operators}\} \\
& \text{qltl}(x, y, (\exists y : \mathbf{G} \text{ trs}) \wedge \mathbf{G} \text{ trs}) \\
\equiv & \{\text{Lemma 6}\} \\
& \text{qltl}(x, y, \mathbf{G} \text{ trs})
\end{aligned}$$

Note that in the above derivation we use the equivalence symbol \equiv , in addition to the equality symbol $=$. Recall that \equiv stands for semantical equivalence of two components (c.f. §4.3.3). On the other hand, $=$ for components means syntactic equality. By definition of the semantics, if two QLTL formulas are equivalent, then the corresponding QLTL components are equivalent.

Based on the above, we define the transformation `stateless2qltl` of a stateless component $C = \text{stateless}(x, y, \text{trs})$, into a QLTL component as follows:

$$\text{stateless2qltl}(C) = \text{qltl}(x, y, \mathbf{G} \text{ trs})$$

5.2 Symbolic Transformations of Special Atomic Components to More General Atomic Components

Based on the lattice in Fig. 5, we define all remaining mappings from more special atomic components to more general atomic components, by composing the previously defined mappings `sts2qltl`, `stateless2qltl`, `stateless2sts`, `det2sts`, `stateless_det2stateless` and `stateless_det2det`, as appropriately.

For mapping stateless deterministic STS components to QLTL components, we have two possibilities: `stateless_det` \rightarrow `det` \rightarrow `sts` \rightarrow `qltl` and `stateless_det` \rightarrow `stateless` \rightarrow `qltl`. We choose the transformation `stateless_det` \rightarrow `stateless` \rightarrow `qltl` because it results in a simpler formula:

$$\text{stateless_det2qltl}(C) = \text{stateless2qltl}(\text{stateless_det2stateless}(C)).$$

Examples. Consider the following STS components:

$$\begin{aligned}
C_1 &= \text{stateless}(x, y, y > x) \\
C_2 &= \text{stateless}(x, (), x > 0) \\
\text{UnitDelay} &= \text{det}(x, s, 0, \text{true}, x, s)
\end{aligned}$$

Then:

$$\begin{aligned}
\text{stateless2qltl}(C_1) &= \text{qltl}(x, y, \mathbf{G} y > x) \\
\text{stateless2qltl}(C_2) &= \text{qltl}(x, (), \mathbf{G} x > 0) \\
\text{det2qltl}(\text{UnitDelay}) &= \text{qltl}(x, y, (\forall s, y : s = 0 \rightarrow \\
&\quad (\text{true} \wedge \bigcirc s = x \wedge y = s) \mathbf{L} (\exists s', y : \text{true} \wedge s' = x \wedge y = s)) \wedge \\
&\quad (\exists s : s = 0 \wedge \mathbf{G} (\text{true} \wedge \bigcirc s = x \wedge y = s)))
\end{aligned}$$

Because $(\exists s', y : \text{true} \wedge s' = x \wedge y = s) \equiv \text{true}$, using Lemma 1 and logical simplifications we obtain:

$$\begin{aligned}
\text{det2qltl(UnitDelay)} &\equiv \text{qltl}(x, y, (\exists s : s = 0 \wedge \mathbf{G}(\bigcirc s = x \wedge y = s))) \\
&\equiv \text{qltl}(x, y, (\exists s : s = 0 \wedge \mathbf{G}(\bigcirc y = x \wedge y = s))) \\
&\equiv \text{qltl}(x, y, (\exists s : s = 0 \wedge \mathbf{G}(y = s) \wedge \mathbf{G}(\bigcirc y = x))) \\
&\equiv \text{qltl}(x, y, (\exists s : y = 0 \wedge \mathbf{G}(y = s) \wedge \mathbf{G}(\bigcirc y = x))) \\
&\equiv \text{qltl}(x, y, (\exists s : \mathbf{G}(y = s)) \wedge y = 0 \wedge \mathbf{G}(\bigcirc y = x)) \\
&\equiv \text{qltl}(x, y, \mathbf{G}(\exists s : y = s) \wedge y = 0 \wedge \mathbf{G}(\bigcirc y = x)) \\
&\equiv \text{qltl}(x, y, y = 0 \wedge \mathbf{G}(\bigcirc y = x))
\end{aligned}$$

5.3 Symbolic Computation of Serial Composition

Given a composite component C formed as the serial composition of two components C_1 and C_2 , i.e., $C = C_1 ; C_2$, we would like to compute a new, *atomic* component C_a , such that C_a is semantically equivalent to C . Because atomic components are represented syntactically (and symbolically) by definition, being able to reduce composite components into atomic components means that we are able to symbolically compute composition operators.

We start by defining the symbolic serial composition of components of the same type.

5.3.1 Symbolic Serial Composition of Two QLTL Components

Let $C = \text{qltl}(x, y, \varphi)$ and $C' = \text{qltl}(y, z, \varphi')$ such that $C ; C'$ is well formed. Then their *symbolic serial composition*, denoted $\text{serial}(C, C')$, is the QLTL component defined by

$$\text{serial}(C, C') = \text{qltl}(x, z, (\forall y : \varphi \rightarrow (\exists z : \varphi')) \wedge (\exists y : \varphi \wedge \varphi')) \quad (12)$$

Note that in the above definition (as well as the ones that follow) we assume that the output variable of C and the input variable of C' have the same name (y) and that the names x , y and z are distinct. In general, this may not be the case. This is not a problem, as we can always rename variables such that this condition is met. Note that variable renaming does not change the semantics of components (c.f. §3.2.2).

The intuition behind the formula in (12) is as follows. The second conjunct $\exists y : \varphi \wedge \varphi'$ ensures that the both contracts φ and φ' of the two components are enforced in the composite contract. The reason we use $\exists y : \varphi \wedge \varphi'$ instead of just the conjunction $\varphi \wedge \varphi'$ is that we want to eliminate (“hide”) internal variable y . (Alternatively, we could also have chosen to output y as an additional output, but would then need an additional hiding operator to remove y .) The first conjunct $\forall y : \varphi \rightarrow (\exists z : \varphi')$ is a formula on the input variable x of the composite component (since all other variables y and z are quantified). This formula restricts the legal inputs of C to those inputs for which, no matter which output C produces, this output is guaranteed to be a legal input for the downstream component C' . For an extensive discussion of the intuition and justification behind this definition, see [77].

5.3.2 Symbolic Serial Composition of Two General STS Components

Let $C = \text{sts}(x, y, s, \text{init}, \text{trs})$ and $C' = \text{sts}(y, z, t, \text{init}', \text{trs}')$ be two general STS components such that $C ; C'$ is well formed. Then:

$$\text{serial}(C, C') = \text{sts}(x, z, (s, t), \text{init} \wedge \text{init}', (\exists s', y : \text{trs}) \wedge (\forall s', y : \text{trs} \rightarrow (\exists t', z : \text{trs}')) \wedge (\exists y : \text{trs} \wedge \text{trs}')) \quad (13)$$

5.3.3 Symbolic Serial Composition of Two Stateless STS Components

Let $C = \text{stateless}(x, y, \text{trs})$ and $C' = \text{stateless}(y, z, \text{trs}')$ be two stateless STS components such that $C ; C'$ is well formed. Then

$$\text{serial}(C, C') = \text{stateless}(x, z, (\forall y : \text{trs} \rightarrow (\exists z : \text{trs}')) \wedge (\exists y : \text{trs} \wedge \text{trs}')) \quad (14)$$

5.3.4 Symbolic Serial Composition of Two Deterministic STS Components

Let $C = \text{det}(x, s, a, p, \text{next}, \text{out})$ and $C' = \text{det}(y, t, b, p', \text{next}', \text{out}')$ be two deterministic STS components such that their serial composition $C ; C'$ is well formed. Then:

$$\text{serial}(C, C') = \text{det}(x, (s, t), (a, b), p \wedge p'[y := \text{out}], (\text{next}, \text{next}'[y := \text{out}]), \text{out}'[y := \text{out}]) \quad (15)$$

where $e[z := e']$ denotes the substitution of all free occurrences of variable z by expression e' in expression e .

5.3.5 Symbolic Serial Composition of Two Stateless Deterministic STS Components

Finally, let $C = \text{stateless_det}(x, p, \text{out})$ and $C' = \text{stateless_det}(y, p', \text{out}')$ be two stateless deterministic STS components such that their serial composition $C ; C'$ is well formed. Then:

$$\text{serial}(C, C') = \text{stateless_det}(x, p \wedge p'[y := \text{out}], \text{out}'[y := \text{out}]) \quad (16)$$

5.3.6 Symbolic Serial Composition of Two Arbitrary Atomic Components

In general, we define the symbolic serial composition of two atomic components C and C' by using the mappings of less general components to more general components (Fig. 5), as appropriate. For example, if C is a deterministic STS component and C' is a stateless STS component, then $\text{serial}(C, C') = \text{serial}(\text{det2sts}(C), \text{stateless2sts}(C'))$. Similarly, if C is a QLTL component and C' is a deterministic STS component, then $\text{serial}(C, C') = \text{serial}(C, \text{det2qltl}(C'))$. Formally, assume that $\text{atm}, \text{atm}' \in \{\text{qltl}, \text{sts}, \text{stateless}, \text{det}, \text{stateless_det}\}$ are the types of the components C and C' , and $\text{common} = \text{atm} \vee \text{atm}'$ is the least general component type that is more general than atm and atm' as defined in Fig. 5. Then

$$\text{serial}(C, C') = \text{serial}(\text{atm2common}(C), \text{atm}'2\text{common}(C')).$$

5.3.7 Correctness of Symbolic Serial Composition

The following theorem demonstrates that our symbolic computations of serial composition are correct, i.e., that the resulting atomic component $\text{serial}(C, C')$ is semantically equivalent to the original composite component $C ; C'$:

Theorem 8. *If C and C' are two atomic components, then*

$$C ; C' \equiv \text{serial}(C, C').$$

Examples. Consider the following STS components:

$$\begin{aligned} C_1 &= \text{stateless}(u, (x, y), \text{true}) \\ C_2 &= \text{stateless_det}((x, y), y \neq 0, \frac{x}{y}) \end{aligned}$$

Then:

$$\begin{aligned} \text{serial}(C_1, C_2) &= \text{serial}(C_1, \text{stateless_det2stateless}(C_2)) \\ &= \text{serial}(\text{stateless}(u, (x, y), \text{true}), \text{stateless}((x, y), z, y \neq 0 \wedge z = \frac{x}{y})) \\ &= \text{stateless}(u, z, (\forall x, y : \text{true} \rightarrow (\exists z : y \neq 0 \wedge z = \frac{x}{y}) \\ &\quad \wedge (\exists x, y : \text{true} \wedge y \neq 0 \wedge z = \frac{x}{y}))) \\ &\equiv \text{stateless}(u, z, \text{false}) \end{aligned}$$

As we can see, the composition results in a stateless STS component with input-output formula `false`. The semantics of such a component is `Fail`, indicating that C_1 and C_2 are *incompatible*. Indeed, in the case of $C_1 ; C_2$, the issue is that C_2 requires its second input, y , to be non-zero, but C_1 cannot guarantee that. The reason is that the input-output formula of C_1 is `true`, meaning that, no matter what its input u is, C_1 may output *any* value for x and y , non-deterministically. This violates the input requirements of C_2 , causing an incompatibility. We will return to this point in §5.8. We also note that this type of incompatibility is impossible to prevent, by controlling the input u . In the example that follows, we see a case where the two components are *not* incompatible, because the input requirements of the downstream component can be met by strengthening the input assumptions of the upstream component:

Consider the following QLTL components:

$$\begin{aligned} C_3 &= \text{qltl}(x, y, \mathbf{G}(x \rightarrow \mathbf{F} y)) \\ C_4 &= \text{qltl}(y, (), \mathbf{G F} y) \end{aligned}$$

Then:

$$\begin{aligned} \text{serial}(C_3, C_4) &= \text{serial}(\text{qltl}(x, y, \mathbf{G}(x \rightarrow \mathbf{F} y)), \text{qltl}(y, (), \mathbf{G F} y)) \\ &= \text{qltl}(x, (), (\forall y : (\mathbf{G}(x \rightarrow \mathbf{F} y)) \rightarrow \mathbf{G F} y) \wedge (\exists y : \mathbf{G}(x \rightarrow \mathbf{F} y) \wedge \mathbf{G F} y)) \\ &\equiv \text{qltl}(x, (), \mathbf{G F} x) \end{aligned}$$

In this example, the downstream component C_4 requires its input y to be infinitely often true ($\mathbf{G F} y$). This can be achieved only if the input x of the upstream component is infinitely often true, which is the condition derived by the serial composition of C_3 and C_4 ($\mathbf{G F} x$). Notice that C_3 does not impose any a priori requirements on its input. However, its input-output relation is the so-called *request-response property* which can be expressed as: *whenever the input x is true, the output y will eventually become true afterwards* ($\mathbf{G}(x \rightarrow \mathbf{F} y)$). This request-response property implies that in order for y to be infinitely-often true, x must be infinitely-often true. Moreover, this is the weakest possible condition that can be enforced on x in order to guarantee that the condition on y holds.

5.4 Symbolic Computation of Parallel Composition

Given a composite component formed as the parallel composition of two components, $C = C_1 \parallel C_2$, we would like to compute an atomic component C_a , such that C_a is semantically equivalent to C . We show how this can be done in this subsection. The symbolic computation of parallel composition follows the same pattern as the one for serial composition (§5.3).

5.4.1 Symbolic Parallel Composition of Two QLTL Components

Let $C = \text{qltl}(x, y, \varphi)$ and $C' = \text{qltl}(u, v, \varphi')$. Then their *symbolic parallel composition*, denoted $\text{parallel}(C, C')$, is the QLTL component defined by

$$\text{parallel}(C, C') = \text{qltl}((x, u), (y, v), \varphi \wedge \varphi') \tag{17}$$

As in the case of symbolic serial composition, we assume that variable names x, y, u, v are all distinct. If this is not the case, then we rename variables as appropriately.

5.4.2 Symbolic Parallel Composition of Two General STS Components

Let $C = \text{sts}(x, y, s, \text{init}, \text{trs})$ and $C' = \text{sts}(u, v, t, \text{init}', \text{trs}')$. Then:

$$\text{parallel}(C, C') = \text{sts}((x, u), (y, v), (s, t), \text{init} \wedge \text{init}', \text{trs} \wedge \text{trs}') \tag{18}$$

5.4.3 Symbolic Parallel Composition of Two Stateless STS Components

Let $C = \text{stateless}(x, y, trs)$ and $C' = \text{stateless}(u, v, trs')$. Then

$$\text{parallel}(C, C') = \text{stateless}((x, u), (y, v), trs \wedge trs') \quad (19)$$

5.4.4 Symbolic Parallel Composition of Two Deterministic STS Components

Let $C = \text{det}(x, s, a, p, next, out)$ and $C' = \text{det}(u, t, b, p', next', out')$. Then:

$$\text{parallel}(C, C') = \text{det}((x, u), (s, t), (a, b), p \wedge p', (next, next'), (out, out')) \quad (20)$$

5.4.5 Symbolic Parallel Composition of Two Stateless Deterministic STS Components

Let $C = \text{stateless_det}(x, p, out)$ and $C' = \text{stateless_det}(u, p', out')$. Then:

$$\text{parallel}(C, C') = \text{stateless_det}((x, u), p \wedge p', (out, out')) \quad (21)$$

5.4.6 Symbolic Parallel Composition of Two Arbitrary Atomic Components

Similar to the symbolic serial composition, we define the symbolic parallel composition of two atomic components C and C' by using the mappings of less general components to more general components (Fig. 5), as appropriate. Formally, assume that $atm, atm' \in \{\text{qltl}, \text{sts}, \text{stateless}, \text{det}, \text{stateless_det}\}$ are the types of the components C and C' , and $common = atm \vee atm'$ is the least general component type that is more general than atm and atm' as defined in Fig. 5. Then

$$\text{parallel}(C, C') = \text{parallel}(atm2common(C), atm'2common(C')).$$

5.4.7 Correctness of Symbolic Parallel Composition

The following theorem demonstrates that our symbolic computations of parallel composition are also correct, i.e., that the resulting atomic component $\text{parallel}(C, C')$ is semantically equivalent to the original composite component $C \parallel C'$:

Theorem 9. *If C and C' are two atomic components, then*

$$C \parallel C' \equiv \text{parallel}(C, C').$$

5.5 Symbolic Computation of Feedback Composition for Decomposable Deterministic STS Components

5.5.1 Decomposable Components

We provide a symbolic closed-form expression for the feedback composition of a deterministic STS component, provided such a component is *decomposable*. Intuitively, decomposability captures the fact that the first output of the component, y_1 , does not depend on its first input, x_1 . This ensures that the feedback composition (which connects y_1 to x_1) does not introduce any circular dependencies.

Definition 21 (Decomposability). *Let C be a deterministic STS component $\text{det}((x_1, \dots, x_n), s, a, p, next, (e_1, \dots, e_m))$ or a stateless deterministic STS component $\text{stateless_det}((x_1, \dots, x_n), p, (e_1, \dots, e_m))$. C is called decomposable if x_1 is not free in e_1 .*

Decomposability is illustrated in Fig. 6(a). The figure shows that expression e_1 depends only on inputs x_2, \dots, x_n .

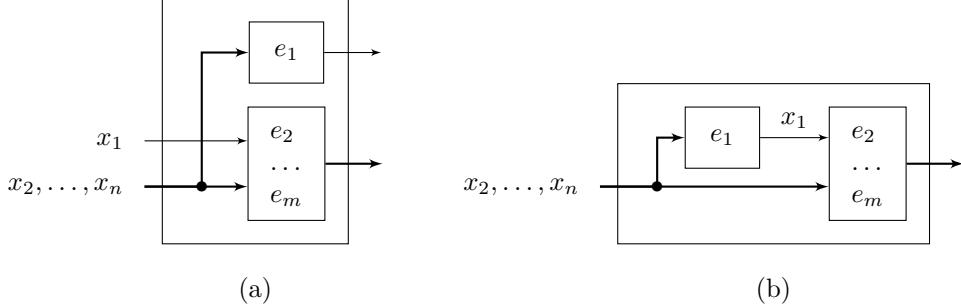


Figure 6: (a) Decomposable deterministic component; (b) the same component after applying feedback, connecting its first output to x_1 .

5.5.2 Symbolic Feedback of a Decomposable Deterministic STS Component

For a decomposable deterministic STS component $C = \text{det}((x_1, \dots, x_n), s, a, p, \text{next}, (e_1, \dots, e_m))$, its *symbolic feedback composition*, denoted $\text{feedback}(C)$, is the deterministic STS component defined by

$$\text{feedback}(C) = \text{det}((x_2, \dots, x_n), s, a, p[x_1 := e_1], \text{next}[x_1 := e_1], (e_2[x_1 := e_1], \dots, e_m[x_1 := e_1])) \quad (22)$$

Thus, computing feedback symbolically consists in removing the first input of the component and replacing the corresponding variable x_1 by the expression of the first output, e_1 , everywhere where x_1 appears. The **feedback** operator is illustrated in Fig. 6(b).

5.5.3 Symbolic Feedback of a Decomposable Stateless Deterministic STS Component

For a decomposable stateless deterministic STS component $C = \text{stateless_det}((x_1, \dots, x_n), p, (e_1, \dots, e_m))$, $\text{feedback}(C)$ is the stateless deterministic STS component defined by

$$\text{feedback}(C) = \text{stateless_det}((x_2, \dots, x_n), p[x_1 := e_1], (e_2[x_1 := e_1], \dots, e_m[x_1 := e_1])) \quad (23)$$

5.5.4 Correctness of Symbolic Feedback Composition

Theorem 10. *If C is a decomposable deterministic STS component, then*

$$\text{feedback}(C) \equiv \text{fdbk}(C)$$

Providing closed-form symbolic computations of feedback composition for general components, including possibly non-deterministic STS and QLTL components, is an open problem, beyond the scope of the current paper. We remark that the straightforward idea of adding to the contract the equality constraint $x = y$ where y is the output connected in feedback to input x , does not work.²

In fact, even obtaining a semantically consistent compositional definition of feedback for non-deterministic and non-input-receptive systems is a challenging problem [65]. Nevertheless, the results that we provide here are sufficient to cover the majority of cases in practice. In particular, the symbolic operator **feedback** can be used to handle Simulink diagrams, provided these diagrams do not contain *algebraic loops*, i.e., circular dependencies (see §5.7).

² One of several problems of this definition is that it does not preserve refinement. For example, the stateless component with contract $x \neq y$ refines the stateless component with contract **true**. Adding the constraint $x = y$ to both contracts yields the components with contracts $x = y$ and **false**, respectively, where the latter no longer refines the former. For a more detailed discussion, see [65].

5.6 Closure Properties of MPT Subclasses w.r.t. Composition Operators

In addition to providing symbolic computation procedures, the results of the above subsections also prove closure properties of the various MPT subclasses of RCRS with respect to the three composition operators. These closure properties are summarized in Tables 1 and 2.

In a nutshell, both serial and parallel composition preserve the most general type of the composed components, according to the lattice in Fig. 5. For instance, the serial (or parallel) composition of two stateless STS components is a stateless STS component; the serial (or parallel) composition of a stateless STS component and a general STS component is a general STS component; and so on. Feedback preserves the type of its component (deterministic or stateless deterministic).

<code>; and </code>	<code>qltl</code>	<code>sts</code>	<code>stateless</code>	<code>det</code>	<code>stateless_det</code>
<code>qltl</code>	<code>qltl</code>	<code>qltl</code>	<code>qltl</code>	<code>qltl</code>	<code>qltl</code>
<code>sts</code>	<code>qltl</code>	<code>sts</code>	<code>sts</code>	<code>sts</code>	<code>sts</code>
<code>stateless</code>	<code>qltl</code>	<code>sts</code>	<code>stateless</code>	<code>sts</code>	<code>stateless</code>
<code>det</code>	<code>qltl</code>	<code>sts</code>	<code>sts</code>	<code>det</code>	<code>det</code>
<code>stateless_det</code>	<code>qltl</code>	<code>sts</code>	<code>stateless</code>	<code>det</code>	<code>stateless_det</code>

Table 1: Closure properties of serial and parallel compositions. The table is to be read as follows: given atomic components C_1, C_2 of types as specified in a row/column pair, the serial or parallel composition of C_1 and C_2 is an atomic component of type as specified in the corresponding table entry.

<code>fbk</code>	<code>det</code> and decomposable	<code>stateless_det</code> and decomposable
	<code>det</code>	<code>stateless_det</code>

Table 2: Closure properties of feedback composition.

5.7 Symbolic Simplification of Arbitrary Composite Components

The results of the previous subsections show how to simplify into an atomic component the serial or parallel composition of two atomic components, or the feedback composition of an atomic component. We can combine these techniques in order to provide a general symbolic simplification algorithm: the algorithm takes as input an arbitrarily complex composite component, and returns an equivalent atomic component. The algorithm is shown in Fig. 7.

The algorithm fails only in case it encounters the feedback of a non-decomposable component. Recall that decomposability implies determinism (c.f. §5.5.1), which means that the test C'' is decomposable means that C'' is of the form `det`((x_1, \dots), $s, a, p, next, (e_1, \dots)$) or `stateless_det`((x_1, \dots), $p, (e_1, \dots)$) and x_1 is not free in e_1 .

We note that in practice, our RCRS implementation on top of Isabelle performs more simplifications in addition to those performed by the procedure **atomic**. For instance, our implementation may be able to simplify a logical formula ϕ into an equivalent but simpler formula ϕ' (e.g., by eliminating quantifiers from ϕ), and consequently also simplify a component, say, `qltl`(x, y, ϕ) into an equivalent but simpler component `qltl`(x, y, ϕ'). These simplifications very much depend on the logic used in the components. Describing the simplifications that our implementation performs is outside the scope of the current paper, as it belongs in the realm of computational logic. It suffices to say that our tool is not optimized for this purpose, and could leverage specialized tools and relevant advances in the field of computational logic.

5.7.1 Deterministic and Algebraic Loop Free Composite Components

In order to state and prove correctness of the algorithm, we extend the notion of determinism to a composite component. We also introduce the notion of *algebraic loop free* components, which capture systems with no

```

atomic( $C$ ) :
  if  $C$  is atomic then
    return  $C$ 
  else if  $C$  is  $C' ; C''$  then
    return serial(atomic( $C'$ ), atomic( $C''$ ))
  else if  $C$  is  $C' \parallel C''$  then
    return parallel(atomic( $C'$ ), atomic( $C''$ ))
  else if  $C$  is fdbk( $C'$ ) then
     $C'' := \text{atomic}(C')$ 
    if  $C''$  is decomposable then
      return feedback( $C''$ )
    else
      fail
  else /* impossible by definition of syntax */
    fail

```

Figure 7: Simplification algorithm for arbitrary composite components.

circular and instantaneous input-output dependencies.

A (possibly composite) component C is said to be *deterministic* if every atomic component of C is either a deterministic STS component or a stateless deterministic STS component. Formally, C is deterministic iff **determ**(C) is true, where **determ**(C) is defined inductively on the structure of C :

$$\begin{aligned}
\mathbf{determ}(\mathbf{det}(x, s, a, p, n, e)) &= \mathbf{true} \\
\mathbf{determ}(\mathbf{stateless_det}(x, p, e)) &= \mathbf{true} \\
\mathbf{determ}(\mathbf{sts}(x, y, s, \mathbf{init}, \mathbf{trs})) &= \mathbf{false} \\
\mathbf{determ}(\mathbf{stateless}(x, y, \mathbf{trs})) &= \mathbf{false} \\
\mathbf{determ}(C ; C') &= \mathbf{determ}(C) \wedge \mathbf{determ}(C') \\
\mathbf{determ}(C \parallel C') &= \mathbf{determ}(C) \wedge \mathbf{determ}(C') \\
\mathbf{determ}(\mathbf{fdbk}(C)) &= \mathbf{determ}(C)
\end{aligned}$$

Notice that this notion of determinism applies to a generally *composite* component C , i.e., a syntactic term in our algebra of components, involving atomic components possibly composed via the three composition operators. This notion of determinism is the generalization of the syntactically deterministic STS components, which are atomic. This notion of determinism is also distinct from any *semantic* notion of determinism (which we have not introduced at all in this paper, as it is not needed).

For a deterministic component we define its *output input dependency relation*. Let C be deterministic, and let $\Sigma_{in}(C) = X_1 \times \dots \times X_n$ and $\Sigma_{out}(C) = Y_1 \times \dots \times Y_m$. The relation $\mathbf{OI}(C) \subseteq \{1, \dots, m\} \times \{1, \dots, n\}$ is defined inductively on the structure of C :

$$\begin{aligned}
\mathbf{OI}(\mathbf{det}((x_1, \dots, x_n), s, a, p, \mathbf{next}, (e_1, \dots, e_m))) &= \{(i, j) \mid x_j \text{ is free in } e_i\} \\
\mathbf{OI}(\mathbf{stateless_det}((x_1, \dots, x_n), p, (e_1, \dots, e_m))) &= \{(i, j) \mid x_j \text{ is free in } e_i\} \\
\mathbf{OI}(C ; C') &= \mathbf{OI}(C') \circ \mathbf{OI}(C) \\
\mathbf{OI}(C \parallel C') &= \mathbf{OI}(C) \cup \{(i + m, j + n) \mid (i, j) \in \mathbf{OI}(C')\} \text{ where} \\
&\quad \Sigma_{in}(C) = X_1 \times \dots \times X_n \wedge \Sigma_{out}(C) = Y_1 \times \dots \times Y_m \\
\mathbf{OI}(\mathbf{fdbk}(C)) &= \{(i, j) \mid i > 0 \wedge j > 0 \wedge ((i + 1, j + 1) \in \mathbf{OI}(C) \\
&\quad \vee ((i + 1, 1) \in \mathbf{OI}(C) \wedge (1, j + 1) \in \mathbf{OI}(C)))\}
\end{aligned}$$

The **OI** relation is preserved by the symbolic operations, as shown by the following lemma:

Lemma 8. If C and C' are deterministic STS components, then

$$\begin{aligned}\mathbf{OI}(C; C') &= \mathbf{OI}(\text{serial}(C, C')) \\ \mathbf{OI}(C \parallel C') &= \mathbf{OI}(\text{parallel}(C, C')).\end{aligned}$$

If C is also decomposable, then

$$\mathbf{OI}(\mathbf{fdbk}(C)) = \mathbf{OI}(\mathbf{feedback}(C)).$$

We introduce now the the notion of algebraic loop free component. Intuitively, a (possibly composite) deterministic component C is algebraic loop free if, whenever C contains a subterm of the form $\mathbf{fdbk}(C')$, the first output of C' does not depend on its first input. This implies that whenever a feedback connection is formed, no circular dependency is introduced. It also ensures that the simplification algorithm will never fail. Formally, for a component C such that $\mathbf{determ}(C)$ is true, $\mathbf{loop-free}(C)$ is defined inductively on the structure of C :

$$\begin{aligned}\mathbf{loop-free}(C) &= \text{true if } C = \text{det}(x, s, a, p, \text{next}, \text{out}) \\ \mathbf{loop-free}(C) &= \text{true if } C = \text{stateless_det}(x, p, \text{out}) \\ \mathbf{loop-free}(C; C') &= \mathbf{loop-free}(C) \wedge \mathbf{loop-free}(C') \\ \mathbf{loop-free}(C \parallel C') &= \mathbf{loop-free}(C) \wedge \mathbf{loop-free}(C') \\ \mathbf{loop-free}(\mathbf{fdbk}(C)) &= \mathbf{loop-free}(C) \wedge (1, 1) \notin \mathbf{OI}(C)\end{aligned}$$

5.7.2 Correctness of the Simplification Algorithm

Theorem 11. Let C be a (possibly composite) component.

1. If C does not contain any \mathbf{fdbk} operators then $\mathbf{atomic}(C)$ does not fail and returns an atomic component such that $\mathbf{atomic}(C) \equiv C$.
2. If $\mathbf{determ}(C) \wedge \mathbf{loop-free}(C)$ is true then $\mathbf{atomic}(C)$ does not fail and returns an atomic component such that $\mathbf{atomic}(C) \equiv C$. Moreover, $\mathbf{atomic}(C)$ is a deterministic STS component and

$$\mathbf{OI}(C) = \mathbf{OI}(\mathbf{atomic}(C)).$$

Proof. The first part of the theorem is a consequence of the fact that the symbolic serial and parallel compositions are defined for all atomic components and return equivalent atomic components, by Theorems 8 and 9.

For the second part, since we have a recursive procedure, we prove its correctness by asumming the correctness of the recursive calls. Additionally, the termination of this procedure is ensured by the fact that all recursive calls are made on “smaller” components. Specifically: we assume that both $\mathbf{determ}(C)$ and $\mathbf{loop-free}(C)$ hold; and we prove that $\mathbf{atomic}(C)$ does not fail, $\mathbf{atomic}(C)$ is a deterministic STS component, $\mathbf{atomic}(C) \equiv C$, and $\mathbf{OI}(C) = \mathbf{OI}(\mathbf{atomic}(C))$.

We only consider the case $C = \mathbf{fdbk}(C')$. All other cases are similar, but simpler. Because $\mathbf{determ}(C)$ and $\mathbf{loop-free}(C)$ hold, we have that $\mathbf{determ}(C')$ and $\mathbf{loop-free}(C')$ also hold, and in addition $(1, 1) \notin \mathbf{OI}(C')$. Using the correctness assumption for the recursive call we have that $\mathbf{atomic}(C')$ does not fail, $C'' = \mathbf{atomic}(C')$ is a deterministic STS component, $C'' \equiv C'$, and $\mathbf{OI}(C'') = \mathbf{OI}(C')$.

Because C'' is a deterministic STS component and $(1, 1) \notin \mathbf{OI}(C') = \mathbf{OI}(C'')$, C'' is decomposable. From this we have that $D := \mathbf{feedback}(C'')$ is defined. Therefore, $\mathbf{atomic}(C)$ returns D and does not fail. It remains to show that D has the desired properties. By the definition of $\mathbf{feedback}(C'')$ and the fact that C'' is a decomposable deterministic STS component, D is also a deterministic STS component. We also have:

$$D = \mathbf{feedback}(C'') \equiv \mathbf{fdbk}(C'') \equiv \mathbf{fdbk}(C') = C$$

where $\mathbf{feedback}(C'') \equiv \mathbf{fdbk}(C'')$ follows from Theorem 10 and $\mathbf{fdbk}(C'') \equiv \mathbf{fdbk}(C')$ follows from $C'' \equiv C'$ and the semantics of \mathbf{fdbk} .

Finally, using Lemma 8 and $\mathbf{OI}(C'') = \mathbf{OI}(C')$, we have

$$\mathbf{OI}(D) = \mathbf{OI}(\mathbf{feedback}(C'')) = \mathbf{OI}(\mathbf{fdbk}(C'')) = \mathbf{OI}(\mathbf{fdbk}(C')) = \mathbf{OI}(C)$$

□

Example. The simplification algorithm applied to the component from Fig. 2 results in

$$\mathbf{atomic}(\mathbf{Sum}) = \mathbf{atomic}(\mathbf{fdbk}(\mathbf{Add}; \mathbf{UnitDelay}; \mathbf{Split})) = \mathbf{det}(y, s, 0, s + y, s).$$

To see how the above is derived, let us first calculate $\mathbf{atomic}(\mathbf{Add}; \mathbf{UnitDelay}; \mathbf{Split})$:

$$\begin{aligned} & \mathbf{atomic}(\mathbf{Add}; \mathbf{UnitDelay}; \mathbf{Split}) \\ = & \mathbf{serial}(\mathbf{atomic}(\mathbf{Add}; \mathbf{UnitDelay}), \mathbf{Split}) \\ = & \mathbf{serial}(\mathbf{serial}(\mathbf{Add}, \mathbf{UnitDelay}), \mathbf{Split}) \\ = & \{\text{Expanding } \mathbf{Add} \text{ and } \mathbf{UnitDelay} \text{ and choosing suitable variable names}\} \\ & \mathbf{serial}(\mathbf{serial}(\mathbf{stateless_det}((x, y), \mathbf{true}, x + y), \mathbf{det}(z, s, 0, \mathbf{true}, z, s)), \mathbf{Split}) \\ = & \{\text{Computing inner } \mathbf{serial}\} \\ & \mathbf{serial}(\mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, s), \mathbf{Split}) \\ = & \{\text{Expanding } \mathbf{Split} \text{ and choosing suitable variable names}\} \\ & \mathbf{serial}(\mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, s), \mathbf{stateless_det}(z, \mathbf{true}, (z, z))) \\ = & \{\text{Computing remaining } \mathbf{serial}\} \\ & \mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, (s, s)) \end{aligned}$$

We calculate now $\mathbf{atomic}(\mathbf{fdbk}(\mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, (s, s))))$:

$$\begin{aligned} & \mathbf{atomic}(\mathbf{fdbk}(\mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, (s, s)))) \\ = & \{x \text{ is not free in } s\} \\ & \mathbf{feedback}(\mathbf{det}((x, y), s, 0, \mathbf{true}, x + y, (s, s))) \\ = & \{\text{Computing } \mathbf{feedback}\} \\ & \mathbf{det}(y, s, 0, \mathbf{true}, s + y, s) \end{aligned}$$

5.8 Checking Validity and Compatibility

Recall the example given in §5.3.7, of the serial composition of components C_1 and C_2 , resulting in a component with input-output relation `false`, implying that $\llbracket C_1 ; C_2 \rrbracket = \text{Fail}$. When this occurs, we say that C_1 and C_2 are *incompatible*. We would like to catch such incompatibilities. This amounts to first simplifying the serial composition $C_1 ; C_2$ into an atomic component C , and then checking whether $\llbracket C \rrbracket = \text{Fail}$.

In general, we say that a component C is *valid* if $\llbracket C \rrbracket \neq \text{Fail}$. Given a component C , we can check whether it is valid, as follows. First, we simplify C to obtain an atomic component $C' = \mathbf{atomic}(C)$. If C' is a QLTL component of the form $\mathbf{qltl}(x, y, \varphi)$ then C' is valid iff φ is satisfiable. The same is true if C' is a stateless STS component of the form $\mathbf{stateless}(x, y, \varphi)$. If C' is a general STS component then we can first transform it into a QLTL component and check satisfiability of the resulting QLTL formula.

Theorem 12. *If C is an atomic component of the form $\mathbf{qltl}(x, y, \varphi)$ or $\mathbf{stateless}(x, y, \varphi)$, then $\llbracket C \rrbracket \neq \text{Fail}$ iff φ is satisfiable.*

As mentioned in the introduction, one of the goals of RCRS is to function as a behavioral type system for reactive system modeling frameworks such as Simulink. In the RCRS setting, type checking consists in checking properties such as compatibility of components, as in the example $C_1 ; C_2$ of §5.3.7. When components are compatible, by computing new (stronger) input preconditions like in the example $C_3 ; C_4$ of §5.3.7, RCRS can be seen as enabling to perform behavioral type inference. Indeed, the new derived condition $\mathbf{G} \mathbf{F} x$ in the above example can be seen as an inferred type of the composite component $C_3 ; C_4$.

We note that the decidability of the satisfiability question for φ depends on the logic used and the domains of the variables in the formula. For instance, although φ can be a QLTL formula, if it restricts the set of constants to `true`, has no functional symbols, and only equality as predicate symbol, then it is equivalent to a QPTL formula,³ for which we can use available techniques [74].

5.9 Checking Input-Receptiveness and Computing Legal Inputs Symbolically

Given a component C , we often want to check whether it is *input-receptive*, i.e., whether $\text{legal}(\llbracket C \rrbracket) = \top$, or equivalently, $\llbracket C \rrbracket(\top) = \top$. More generally, we may want to compute the legal input values for C , which is akin to type inference as discussed above. To do this, we will provide a symbolic method to compute $\text{legal}(\llbracket C \rrbracket)$ as a formula $\text{legal}(C)$. Then, checking that C is input-receptive amounts to checking that the formula $\text{legal}(C)$ is valid, or equivalently, checking that $\neg \text{legal}(C)$ is unsatisfiable. We assume that C is atomic (otherwise, we first simplify C using the algorithm of §5.7).

Definition 22. Given an atomic component C , we define $\text{legal}(C)$, a formula characterizing the legal inputs of C . $\text{legal}(C)$ is defined based on the type of C :

$$\text{legal}(\text{qltl}(x, y, \varphi)) = (\exists y : \varphi) \quad (24)$$

$$\text{legal}(\text{sts}(x, y, s, \text{init}, \text{trs})) = (\forall s, y : \text{init} \rightarrow (\varphi \mathbf{L} \varphi')) \quad (25)$$

$$\text{legal}(\text{stateless}(x, y, \text{trs})) = \mathbf{G}(\exists y : \text{trs}) \quad (26)$$

$$\text{For a deterministic STS component } C, \text{legal}(C) = \text{legal}(\text{det2sts}(C)) \quad (27)$$

$$\text{For a stateless deterministic STS component } C, \text{legal}(C) = \text{legal}(\text{stateless_det2stateless}(C)) \quad (28)$$

where $\varphi = \text{trs}[s' := \bigcirc s]$ and $\varphi' = (\exists s', y : \text{trs})$.

The next theorem shows that **legal** correctly characterizes the semantic predicate **legal**:

Theorem 13. If C is an atomic component , then

$$\text{legal}(\llbracket C \rrbracket) = \{\sigma_x \mid \sigma_x \models \text{legal}(C)\}.$$

It follows from Theorem 13 that a component C is input-receptive iff the formula $\text{legal}(C)$ is valid.

5.10 Checking Refinement Symbolically

We end this section by showing how to check whether a component refines another component. Again, we will assume that the components in question are atomic (if not, they can be simplified using the **atomic** procedure).

Theorem 14. Let $C_1 = \text{sts}(x, y, s, \text{init}, r_1)$, $C'_1 = \text{sts}(x, y, s, \text{init}', r'_1)$, $C_2 = \text{stateless}(x, y, r_2)$, $C'_2 = \text{stateless}(x, y, r'_2)$, $C_3 = \text{qltl}(x, y, \varphi)$, and $C'_3 = \text{qltl}(x, y, \varphi')$. Then:

1. C_1 is refined by C'_1 if the formula

$$(\text{init}' \rightarrow \text{init}) \wedge ((\exists s', y : r_1) \rightarrow (\exists s', y : r'_1)) \wedge ((\exists s', y : r_1) \wedge r'_1 \rightarrow r)$$

is valid.

³ For example, the atomic QLTL formula $\bigcirc \bigcirc x = \bigcirc y$ can be translated into the LTL formula $\mathbf{X} \mathbf{X} x \Leftrightarrow \mathbf{X} y$, and the formula $\bigcirc \bigcirc \bigcirc x = \text{true}$ into $\mathbf{X} \mathbf{X} \mathbf{X} x$.

2. C_2 is refined by C'_2 if and only if the formula

$$((\exists y : r_2) \rightarrow (\exists y : r'_2)) \wedge ((\exists y : r_2) \wedge r'_2 \rightarrow r_2)$$

is valid.

3. C_3 is refined by C'_3 if and only if the formula

$$((\exists y : \varphi) \rightarrow (\exists y : \varphi')) \wedge (((\exists y : \varphi) \wedge \varphi') \rightarrow \varphi)$$

is valid.

As the above theorem shows, checking refinement amounts to checking validity (or equivalently, satisfiability of the negation) of first-order formulas formed by the various symbolic expressions in the component specifications. The exact logic of these formulas depends on the logics used by the components. For example, if C_3 and C'_3 both use quantifier-free LTL for ϕ and ϕ' , then in order to check refinement we need to check satisfiability of a first-order QLTL formula.

Examples. Consider again the QLTL component $C = \text{qltl}(\text{(), } t, \text{oven})$, introduced in §3.3, where

$$\begin{aligned}\text{oven} &= (t = 20 \wedge ((t < \circ t \wedge t < 180) \mathbf{U} \text{thermostat})) \\ \text{thermostat} &= \mathbf{G}(180 \leq t \wedge t \leq 220)\end{aligned}$$

Let us introduce a refined version C' of C :

$$\begin{aligned}C' &= \text{sts}(\text{(), } t, (s, sw), \text{init}, \text{trs}) \text{ where} \\ \text{init} &= s = 20 \wedge sw = \text{on} \\ \text{trs} &= (t = s) \wedge \\ &\quad (\text{if } sw = \text{on} \text{ then } s < s' < s + 5 \text{ else } (\text{if } s > 10 \text{ then } s - 5 < s' < s \text{ else } s' = s)) \wedge \\ &\quad (\text{if } sw = \text{on} \wedge s > 210 \text{ then } sw' = \text{off} \text{ else} \\ &\quad \quad (\text{if } sw = \text{off} \wedge s < 190 \text{ then } sw' = \text{on} \text{ else } sw' = sw))\end{aligned}$$

C' is an STS component with no input variables, output variable t , and state variables s and sw , recording the current temperature of the oven, and the on/off status of the switch, respectively. When sw is **on**, the temperature increases nondeterministically by up to 5 units, otherwise the temperature decreases nondeterministically by up to 5 units. When the temperature exceeds 210, the switch is turned off; when the temperature is below 190, the switch is turned **on**; otherwise sw remains unchanged. The output t is always equal to the current state s . Initially the temperature is 20, and sw is **on**.

Using Theorem 14, and the properties of `sts2qltl` we have:

$$\begin{aligned}C &\preceq C' \\ \Leftrightarrow C &\preceq \text{sts2qltl}(C') \\ \Leftrightarrow q\text{ltl}(\text{(), } t, \text{oven}) &\preceq q\text{ltl}(\text{(), } t, (\forall s, sw, t : \text{init} \rightarrow (\varphi \mathbf{L} \varphi')) \wedge (\exists s, sw : \text{init} \wedge \mathbf{G} \varphi)) \\ &\quad \text{where } \varphi = \text{trs}[s', sw' := \circ s, \circ sw] \text{ and } \varphi' = (\exists s', sw', t : \text{trs}) \\ \Leftrightarrow &\{\text{Using Lemma 1, because } \varphi' \equiv \text{true}\} \\ q\text{ltl}(\text{(), } t, \text{oven}) &\preceq q\text{ltl}(\text{(), } t, (\exists s, sw : \text{init} \wedge \mathbf{G} \varphi)) \\ \Leftrightarrow &\{\text{Using Theorem 14}\} \\ ((\exists t : \text{oven}) \rightarrow (\exists t, s, sw : \text{init} \wedge \mathbf{G} \varphi)) \wedge &((\exists t : \text{oven}) \wedge (\exists s, sw : \text{init} \wedge \mathbf{G} \varphi)) \rightarrow \text{oven} \text{ is valid} \\ \Leftrightarrow &\{\text{Because } (\exists t : \text{oven}) \equiv \text{true} \text{ and } (\exists t, s, sw : \text{init} \wedge \mathbf{G} \varphi) \equiv \text{true}\} \\ ((\exists s, sw : \text{init} \wedge \mathbf{G} \varphi) \rightarrow \text{oven}) &\text{ is valid}\end{aligned}$$

Thus, checking whether C' refines C amounts to checking whether the QLTL formula $((\exists s, sw : \mathbf{init} \wedge \mathbf{G} \varphi) \rightarrow \mathbf{oven})$ is valid. This indeed holds for this example and can be shown using logical reasoning.

The above example is relatively simple in the sense that in the end refinement reduces to checking implication between the corresponding contracts. Indeed, this is always the case for input-receptive systems, as in the example above. However, refinement is *not* equivalent to implication in the general case of non-input-receptive systems. For example, consider the components:

$$\begin{aligned} C_1 &= \mathbf{stateless}(x, y, x \geq 0 \wedge y \geq x) \\ C_2 &= \mathbf{stateless}(x, y, x \leq y \leq x + 10) \end{aligned}$$

Using Theorem 14, we have:

$$\begin{aligned} C_1 &\preceq C_2 \\ \Leftrightarrow &\{\text{Theorem 14}\} \\ &((\exists y : x \geq 0 \wedge y \geq x) \rightarrow (\exists y : x \leq y \leq x + 10)) \wedge \\ &((\exists y : x \geq 0 \wedge y \geq x) \wedge x \leq y \leq x + 10 \rightarrow x \geq 0 \wedge y \geq x) \text{ is valid} \\ \Leftrightarrow &\{\text{Arithmetic and logical reasoning}\} \\ &\mathbf{true} \end{aligned}$$

Note that the second and third parts of Theorem 14 provide necessary *and* sufficient conditions, while the first part only provides a sufficient, but generally not necessary condition. Indeed, the condition is generally not necessary in the case of STS components with state, as state space computation is ignored by the condition. This can be remedied by transforming STS components into equivalent QLTL components and then applying the second part of the theorem. An alternative which may be more tractable, particularly in the case of finite-state systems, is to use techniques akin to strategy synthesis in games, such as those proposed in [77] for finite-state relational interfaces.

Another limitation of the first part of Theorem 14 is that it requires the two STS components to have the same state space, i.e., the same state variable s . This restriction can be lifted using the well-known idea of data refinement [41, 8].

Theorem 15. *Let $C_1 = \mathbf{sts}(x, y, s, init, r)$, $C_2 = \mathbf{sts}(x, y, t, init', r')$ be two STS components, and D a (data refinement) expression on variables s and t . Let $p = (\exists s', y : r)$ and $p' = (\exists t', y : r')$. If the formulas*

$$\begin{aligned} &(\forall t : init' \rightarrow (\exists s : D \wedge init)) \\ &(\forall t, x, s : D \wedge p \rightarrow p') \\ &(\forall t, x, s, t', y : D \wedge p \wedge r' \rightarrow (\exists s' : D[t, s := t', s'] \wedge r)) \end{aligned}$$

are valid, then C_1 is refined by C_2 .

6 Toolset and Case Studies

The RCRS framework comes with a publicly available toolset, which can be downloaded from rcrs.cs.aalto.fi. The toolset, illustrated in Fig. 8, consists of:

- A full implementation of the RCRS theory in the Isabelle theorem prover [58]. The RCRS theory consists of a total of 15100 lines of Isabelle code – 12200 lines for the RCRS formal framework, 1800 lines for the simplification algorithm and 1100 lines for the Simulink component library. The latter library is a characterization of basic Simulink blocks as RCRS components (at the time of writing this paper, 48 of Simulink’s basic blocks can be handled).
- A set of auxiliary tools, in particular, a **Translator** of Simulink diagrams into RCRS code.

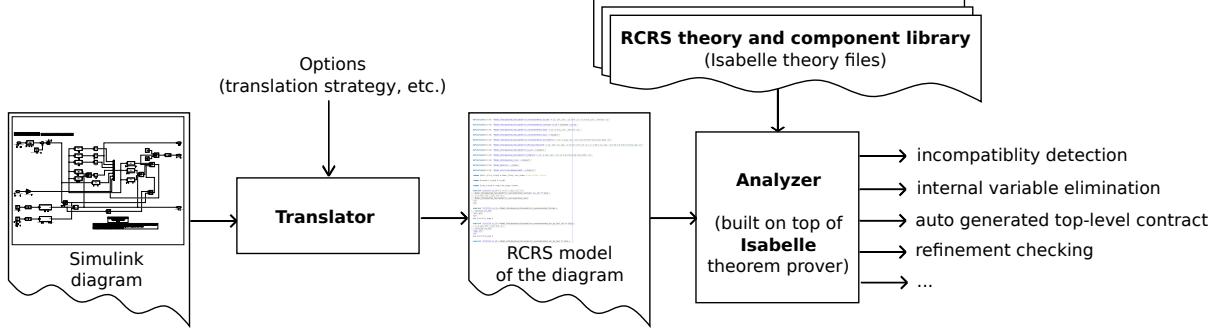


Figure 8: The RCRS toolset.

The Isabelle implementation consists of theories for the various elements of RCRS, including QLTL, the theory of MPTs, etc. We implemented in Isabelle a *shallow embedding* [14] of the language introduced in §3. The advantage of a shallow embedding is that all datatypes of Isabelle are available for specification of components, and we can use the existing Isabelle mechanism for renaming bound variables in compositions. The disadvantage of this shallow embedding is that we cannot express the Algorithm 7 within Isabelle (hence the “manual” proof that we provide for Theorem 11). A *deep embedding*, in which the syntax of components is defined as a datatype of Isabelle, is possible, and is left as an open future work direction.

We implemented Algorithm 7 in SML, the meta-language of Isabelle. The SML program takes as input a component C and returns not only a simplified atomic component $\text{atomic}(C)$, but also a proved Isabelle theorem of the fact $C \equiv \text{atomic}(C)$. The simplification program, as well as a number of other procedures to perform compatibility checking, validity checking, etc., form what we call the **Analyzer** in Fig. 8.

The **Translator** is a Python program that takes as input a Simulink diagram and translates it into RCRS. Specifically, the tool produces an Isabelle theory file which contains: (1) the definition of the composite components representing the Simulink diagram; and (2) embedded bottom-up simplification procedures and the corresponding correctness theorems. By running this theory file in Isabelle, we obtain an atomic component corresponding to the top-level Simulink model, equivalent to the original composite component. As a special case, if the Simulink diagram contains inconsistencies (e.g., division by zero), these are detected by obtaining `Fail` as the top-level atomic component. The error can be localized by finding earlier points in the Simulink hierarchy (subsystems) which already resulted in `Fail`.

As mentioned earlier, how to obtain a composite component from a graphical block diagram is an interesting problem. This problem is studied in depth in [31], where several translation strategies are proposed. These various strategies all yield semantically equivalent components, but with different trade-offs in terms of size, readability, effectiveness of the simplification procedures, and so on. The **Translator** implements all these translation strategies, allowing the user to explore these trade-offs. Further details on the translation problem are provided in [31, 63]. A proof that the translation strategies yield semantically equivalent components is provided in [62].

We have used the RCRS toolset on several case studies, including a real-life benchmark provided by the Toyota motor company. The benchmark involves a Fuel Control System (FCS) described in [42, 43]. FCS aims at controlling the air mass and injected fuel in a car engine such that their ratio is always optimal. This problem has important implications on lowering pollution and costs by improving the engine performance. Toyota has made several versions of FCS publicly available as Simulink models at <https://cps-vo.org/group/ARCH/benchmarks>.

We have used the RCRS toolset to process two of the three Simulink models in the FCS benchmark suite (the third model contains blocks that are currently not implemented in the RCRS component library). A typical model in this set has a 3-layer hierarchy with a total of 104 Simulink block instances (97 basic blocks and 7 subsystems), and 101 connections out of which 8 are feedbacks. Each basic Simulink block is modeled in our framework by an atomic STS component (possibly stateless). These atomic STS component are created once, and form part of the RCRS implementation, which is reused for different Simulink models. The

particular FCS diagram is translated into RCRS using the **Translator**, and simplified within Isabelle using our SML simplification procedure. After simplification, we obtain an atomic deterministic STS component with no inputs, 7 outputs, and 14 state variables. Its contract (which is 8337 characters long) includes a condition on the state variables, in particular, that a certain state variable must always be non-negative (as its value is fed into a square-root block). This condition makes it not immediately obvious that the whole system is valid (i.e., not `Fail`). However, we can show after applying the transformation `sts2qltl` that the resulting formula is satisfiable, which implies that the original model is consistent (i.e., no connections result in incompatibilities, there are no divisions by zero, etc.). This illustrates the use of RCRS as a powerful static analysis tool. More details on the FCS case study are provided in [31, 63].

7 Related Work

Several formal compositional frameworks exist in the literature. Most closely related to RCRS are the frameworks of FOCUS [15], input-output automata [51], reactive modules [6], interface automata [22], and Dill’s trace theory [29]. RCRS shares with these frameworks many key compositionality principles, such as the notion of refinement. At the same time, RCRS differs and complements these frameworks in important ways. Specifically, FOCUS, IO-automata, and reactive modules, are limited to input-receptive systems, while RCRS is explicitly designed to handle non-input-receptive specifications. The benefits of non-input-receptiveness are discussed extensively in [77] and will not be repeated here. Interface automata are a low-level formalism whereas RCRS specifications and reasoning are symbolic. For instance, in RCRS one can naturally express systems with infinite state-spaces, input-spaces, or output-spaces. (Such systems can even be handled automatically, provided the corresponding logic they are expressed in is decidable.) Both interface automata and Dill’s trace theory use a single form of asynchronous parallel composition, whereas RCRS has three primitive composition operators (serial, parallel, feedback) with synchronous semantics.

Our work adapts and extends to the reactive system setting many of the ideas developed previously in a long line of research on correctness and compositionality for sequential programs. This line of research goes back to the works of Floyd, Hoare, Dijkstra, and Wirth, on formal program semantics, weakest preconditions, program development by stepwise refinement, and so on [32, 40, 27, 79]. It also goes back to game-theoretic semantics of sequential programs as developed in the original refinement calculus [10], as well as to contract-based design [54]. Many of the concepts used in our work are in spirit similar to those used in the above works. For instance, an input-output formula ϕ used in an atomic component in our language can be seen as a *contract* between the environment of the component and the component itself: the environment must satisfy the contract by providing to the component legal inputs, and the component must in turn provide legal outputs (for those inputs). On the other hand, several of the concepts used here come from the world of reactive systems and as such do not have a direct correspondence in the world of sequential programs. For instance, this is the case with feedback composition.

RCRS extends refinement calculus from predicate to property transformers. Extensions of refinement calculus to infinite behaviors have also been proposed in the frameworks of *action systems* [11], *fair action systems* [12], and *Event B* [3]. These frameworks use predicate (not property) transformers as semantic foundation; they can handle certain property patterns (e.g., fairness) by providing proof rules for these properties, but they do not treat liveness and LTL properties in general [12, 80, 39]. The *Temporal Logic of Actions* [46] can be used to specify liveness properties, but does not distinguish between inputs and outputs, and as such cannot express non-input-receptive components.

Our specifications can be seen as “rich”, behavioral types [47, 22]. Indeed, our work is closely related to programming languages and specifically type theory. Early works on behavioral types include [57, 48, 26]. Most closely related to our framework is the framework of *liquid types* [67].

Behavioral type frameworks have also been proposed in reactive system settings. In the SimCheck framework, Simulink blocks are annotated with constraints on input and output variables, much like stateless components in RCRS [68]. RCRS is more general as it also allows to specify stateful components. RCRS is also a more complete compositional framework, with composition operators and refinement, which are not considered in [68]. Other behavioral type theories for reactive systems have been proposed in [23, 18, 30].

Compared to RCRS, these works are less general. In particular, [23, 30] are limited to specifications which separate the assumptions on the inputs from the guarantees on the outputs, and as such cannot capture input-output relations. [18] considers a synchronous model which allows to specify legal values of inputs and outputs at the *next* step, given the current state. This model does not allow to capture relations between inputs and outputs within the same step, which RCRS allows.

Our work is related to formal verification frameworks for *hybrid systems* [5]. Broadly speaking, these can be classified into frameworks following a model-checking approach, which typically use automata-based specification languages and state-space exploration techniques, and those following a theorem-proving approach, which typically use logic-based specifications. More closely related to RCRS are the latter, among which, CircusTime [17], KeYmaera [34], and the PVS-based approach in [2]. CircusTime can handle a larger class of Simulink diagrams than the current implementation of RCRS. In particular, CircusTime can handle *multi-rate* diagrams, where different parts of the model work at different rates (periods). On the other hand, CircusTime is based on predicate (not property) transformers, and as such cannot handle liveness properties. KeYmaera is a theorem prover based on *differential dynamic logic* [59], which is itself based on *dynamic logic* [36]. The focus of both KeYmaera and the work in [2] is verification, and not compositionality. For instance, these works do not distinguish between inputs and outputs and do not investigate considerations such as input-receptiveness. The work of [66] distinguishes inputs and outputs, but provides a system model where the output relation is separated from the transition relation, and where the output relation is assumed to be total, meaning that there exists an output for every input and current state combination. This does not allow to specify non-input-receptive stateless components, such as for example the `Div` component from Section 3.

Our component algebra is similar to the algebra of flownomials [75] and to the relational model for non-deterministic dataflow [38]. In [20], graphs and graph operations which can be viewed as block diagrams are represented by algebraic expressions and operations, and a complete equational axiomatization of the equivalence of the graph expressions is given. This is then applied to flow-charts as investigated in [70]. The translation of block diagrams in general and Simulink in particular has been treated in a large number of papers, with various goals, including verification and code generation (e.g., see [78, 53, 19, 49, 71, 13, 81, 82, 83, 55]). Although we share several of the ideas of the above works, our main goal here is not to formalize the language of block diagrams, neither their translation to other formalisms, but to provide a complete compositional framework for reasoning about reactive systems.

RCRS is naturally related to *compositional verification* frameworks, such as [35, 1, 52, 72, 37, 24, 25]. In particular, compositional verification frameworks often make use of a refinement relation such as trace inclusion or simulation [72, 37]. However, the focus of these frameworks is different than that of RCRS. In compositional verification, the focus is to “break down” a large (and usually computationally expensive) verification task into smaller (and hopefully easier to calculate) subtasks. For this purpose, compositional verification frameworks employ several kinds of *decomposition rules*. An example of such a rule is the so-called *precongruence* rule (i.e., preservation of refinement by composition): if P_1 refines Q_1 , and P_2 refines Q_2 , then the composition $P_1 \parallel P_2$ refines $Q_1 \parallel Q_2$. This, together with preservation of properties by refinement, allows us to conclude that $P_1 \parallel P_2$ satisfies some property ϕ , provided we can prove that $Q_1 \parallel Q_2$ satisfies ϕ . The latter might be a simpler verification task, if Q_1 and Q_2 are smaller than P_1 and P_2 . The essence of compositional verification is in finding such *abstract* versions Q_1 and Q_2 of the concrete processes in question, P_1 and P_2 , and employing decomposition rules like the one above in the hope of making verification simpler. RCRS can also be used for compositional verification: indeed, RCRS provides both the precongruence rule, and preservation of properties by refinement. Note that, in traditional settings, precongruence is not always powerful enough, and for this reason most compositional verification frameworks employ more complex decomposition rules (e.g., see [56]). In settings which allow non-input-receptive components, such as ours, there are indications that the precongruence rule is sufficient for compositional verification purposes [73], although more work is required to establish this in the specific context of RCRS. Such work is beyond the scope of the current paper. We also note that, beyond compositional verification with precongruence, RCRS provides a behavioral type theory which allows to state system properties such as compatibility, which is typically not available in compositional verification frameworks.

Refinement can be seen as the inverse of *abstraction*, and as such our framework is related to general frameworks such as *abstract interpretation* [21]. Several abstractions have been proposed in reactive system settings, including *relational* abstractions for hybrid systems, which are related to Simulink [69]. The focus of these works is verification, and abstraction is used as a mechanism to remove details from the model that make verification harder. In RCRS, the simplification procedure that we employ can be seen as an abstraction process, as it eliminates internal variable information. However, RCRS simplification is an *exact* abstraction, in the sense that it does not lose any information: the final system is equivalent to the original one, and not an over- or under-approximation, as is usually the case with typical abstractions for verification purposes.

8 Conclusion

We presented RCRS, a compositional framework for modeling and reasoning about reactive systems. This paper focuses on the theory and methodology of RCRS, its formal semantics, and associated techniques for symbolic and computer-aided reasoning.

RCRS is an ongoing project, and a number of problems remain open. Future work directions include:

- An extension of the framework to systems with algebraic loops, which necessitates handling instantaneous feedback. Here, the preliminary ideas of [65] can be helpful in defining the semantics of instantaneous feedback. However, [65] does not provide solutions on how to obtain symbolic closed-form expression for the feedback of general components.
- Extension of the results of §5.5 to general components, possibly non-deterministic or non-decomposable.
- An extension of the framework to *acausal* systems, i.e., systems without a clear distinction of inputs and outputs [33].
- The development of better symbolic reasoning techniques, such as simplification of logical formulas, decision procedures, etc.

References

- [1] M. Abadi and L. Lamport. Conjoining specifications. *ACM Trans. Program. Lang. Syst.*, 17(3):507–535, 1995.
- [2] E. Ábrahám-Mumm, M. Steffen, and U. Hannemann. Verification of hybrid systems: Formalization and proof rules in PVS. In *7th Intl. Conf. Engineering of Complex Computer Systems (ICECCS 2001)*, pages 48–57, 2001.
- [3] J.-R. Abrial. *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, New York, NY, USA, 1st edition, 2010.
- [4] B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181 – 185, 1985.
- [5] R. Alur, C. Courcoubetis, N. Halbwachs, T. Henzinger, P. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.
- [6] R. Alur and T. Henzinger. Reactive modules. *Formal Methods in System Design*, 15:7–48, 1999.
- [7] R. Alur, T. Henzinger, O. Kupferman, and M. Vardi. Alternating refinement relations. In *CONCUR’98*, volume 1466 of *LNCS*. Springer, 1998.
- [8] R. J. Back. *Correctness preserving program refinements: proof theory and applications*, volume 131 of *Mathematical Centre Tracts*. Mathematisch Centrum, Amsterdam, 1980.

- [9] R.-J. Back and M. Butler. *Exploring summation and product operators in the refinement calculus*, pages 128–158. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995.
- [10] R.-J. Back and J. von Wright. *Refinement Calculus: A Systematic Introduction*. Springer, 1998.
- [11] R.-J. Back and J. Wright. Trace refinement of action systems. In B. Jonsson and J. Parrow, editors, *CONCUR '94: Concurrency Theory*, volume 836 of *Lecture Notes in Computer Science*, pages 367–384. Springer Berlin Heidelberg, 1994.
- [12] R.-J. Back and Q. Xu. Refinement of fair action systems. *Acta Informatica*, 35(2):131–165, 1998.
- [13] P. Boström. Contract-Based Verification of Simulink Models. In S. Qin and Z. Qiu, editors, *Formal Methods and Software Engineering*, volume 6991 of *Lecture Notes in Computer Science*, pages 291–306. Springer Berlin Heidelberg, 2011.
- [14] R. J. Boulton, A. Gordon, M. J. C. Gordon, J. Harrison, J. Herbert, and J. V. Tassel. Experience with embedding hardware description languages in HOL. In *IFIP TC10/WG 10.2 Intl. Conf. on Theorem Provers in Circuit Design*, pages 129–156. North-Holland Publishing Co., 1992.
- [15] M. Broy and K. Stølen. *Specification and development of interactive systems: focus on streams, interfaces, and refinement*. Springer, 2001.
- [16] R. Cavada, A. Cimatti, M. Dorigatti, A. Griggio, A. Mariotti, A. Micheli, S. Mover, M. Roveri, and S. Tonetta. The nuxmv symbolic model checker. In A. Biere and R. Bloem, editors, *Computer Aided Verification: 26th International Conference, CAV 2014*, pages 334–342, Cham, 2014. Springer.
- [17] A. L. C. Cavalcanti, A. Mota, and J. C. P. Woodcock. Simulink timed models for program verification. In Z. Liu, J. C. P. Woodcock, and H. Zhu, editors, *Theories of Programming and Formal Methods - Essays Dedicated to Jifeng He on the Occasion of His 70th Birthday*, volume 8051 of *Lecture Notes in Computer Science*, pages 82–99. Springer, 2013.
- [18] A. Chakrabarti, L. de Alfaro, T. Henzinger, and F. Mang. Synchronous and bidirectional component interfaces. In *CAV*, LNCS 2404, pages 414–427. Springer, 2002.
- [19] C. Chen, J. S. Dong, and J. Sun. A formal framework for modeling and validating Simulink diagrams. *Formal Aspects of Computing*, 21(5):451–483, 2009.
- [20] B. Courcelle. *A representation of graphs by algebraic expressions and its use for graph rewriting systems*, pages 112–132. Springer Berlin Heidelberg, Berlin, Heidelberg, 1987.
- [21] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *4th ACM Symp. POPL*, 1977.
- [22] L. de Alfaro and T. Henzinger. Interface automata. In *Foundations of Software Engineering (FSE)*. ACM Press, 2001.
- [23] L. de Alfaro and T. Henzinger. Interface theories for component-based design. In *EMSOFT'01*. Springer, LNCS 2211, 2001.
- [24] W. de Roever, H. Langmaack, and A. E. Pnueli. *Compositionality: The Significant Difference*. LNCS. Springer, 1998.
- [25] W.-P. de Roever, F. de Boer, U. Hanneman, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Non-compositional Methods*. Cambridge University Press, 2012.
- [26] K. Dhara and G. Leavens. Forcing behavioral subtyping through specification inheritance. In *ICSE'96: 18th Intl. Conf. on Software Engineering*, pages 258–267. IEEE Computer Society, 1996.

- [27] E. Dijkstra. Notes on structured programming. In O. Dahl, E. Dijkstra, and C. Hoare, editors, *Structured programming*, pages 1–82. Academic Press, London, UK, 1972.
- [28] E. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Comm. ACM*, 18(8):453–457, 1975.
- [29] D. Dill. *Trace theory for automatic hierarchical verification of speed-independent circuits*. MIT Press, Cambridge, MA, USA, 1987.
- [30] L. Doyen, T. Henzinger, B. Jobstmann, and T. Petrov. Interface theories with component reuse. In *8th ACM & IEEE International conference on Embedded software, EMSOFT*, pages 79–88, 2008.
- [31] I. Dragomir, V. Preoteasa, and S. Tripakis. Compositional Semantics and Analysis of Hierarchical Block Diagrams. In *SPIN*, pages 38–56. Springer, 2016.
- [32] R. Floyd. Assigning meanings to programs. In *In. Proc. Symp. on Appl. Math. 19*, pages 19–32. American Mathematical Society, 1967.
- [33] P. Fritzson. *Principles of Object-Oriented Modeling and Simulation with Modelica 3.3: A Cyber-Physical Approach*. Wiley, 2 edition, 2014.
- [34] N. Fulton, S. Mitsch, J.-D. Quesel, M. Völp, and A. Platzer. KeYmaera X: An axiomatic tactical theorem prover for hybrid systems. In A. P. Felty and A. Middeldorp, editors, *CADE*, volume 9195 of *LNCS*, pages 527–538. Springer, 2015.
- [35] O. Grumberg and D. Long. Model checking and modular verification. *ACM Trans. Program. Lang. Syst.*, 16(3):843–871, 1994.
- [36] D. Harel, J. Tiuryn, and D. Kozen. *Dynamic Logic*. MIT Press, 2000.
- [37] T. Henzinger, S. Qadeer, and S. Rajamani. You assume, we guarantee: Methodology and case studies. In *CAV’98*, volume 1427 of *LNCS*. Springer-Verlag, 1998.
- [38] T. T. Hildebrandt, P. Panangaden, and G. Winskel. A relational model of non-deterministic dataflow. *Mathematical Structures in Computer Science*, 14(5):613–649, 10 2004.
- [39] T. S. Hoang and J.-R. Abrial. Reasoning about liveness properties in event-b. In *Proceedings of the 13th International Conference on Formal Methods and Software Engineering, ICFEM’11*, pages 456–471, Berlin, Heidelberg, 2011. Springer-Verlag.
- [40] C. A. R. Hoare. An axiomatic basis for computer programming. *Comm. ACM*, 12(10):576–580, 1969.
- [41] C. A. R. Hoare. Proof of correctness of data representations. *Acta Informatica*, 1(4), Dec. 1972.
- [42] X. Jin, J. Deshmukh, J. Kapinski, K. Ueda, and K. Butts. Benchmarks for model transformations and conformance checking. In *1st Intl. Workshop on Applied Verification for Continuous and Hybrid Systems (ARCH)*, 2014.
- [43] X. Jin, J. V. Deshmukh, J. Kapinski, K. Ueda, and K. Butts. Powertrain Control Verification Benchmark. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control, HSCC’14*, pages 253–262. ACM, 2014.
- [44] Y. Kesten, Z. Manna, and A. Pnueli. Temporal verification of simulation and refinement. In J. W. de Bakker, W. P. de Roever, and G. Rozenberg, editors, *A Decade of Concurrency Reflections and Perspectives: REX School/Symposium*, pages 273–346. Springer, 1994.
- [45] Y. Kesten and A. Pnueli. Complete proof system for qptl. *Journal of Logic and Computation*, 12(5):701, 2002.

- [46] L. Lamport. The temporal logic of actions. *ACM Trans. Program. Lang. Syst.*, 16(3):872–923, May 1994.
- [47] E. Lee and Y. Xiong. System-level types for component-based design. In *EMSOFT’01: 1st Intl. Workshop on Embedded Software*, pages 237–253. Springer, 2001.
- [48] B. Liskov and J. Wing. A behavioral notion of subtyping. *ACM Trans. Program. Lang. Syst.*, 16(6):1811–1841, 1994.
- [49] R. Lublinerman, C. Szegedy, and S. Tripakis. Modular code generation from synchronous block diagrams – modularity vs. code size. In *36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL’09)*, pages 78–89. ACM, Jan. 2009.
- [50] R. Lublinerman and S. Tripakis. Modularity vs. reusability: Code generation from synchronous block diagrams. In *Design, Automation, and Test in Europe (DATE’08)*, pages 1504–1509. ACM, Mar. 2008.
- [51] N. Lynch and M. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2:219–246, 1989.
- [52] K. McMillan. A compositional rule for hardware design refinement. In *Computer Aided Verification (CAV’97)*, volume 1254 of *LNCS*. Springer-Verlag, 1997.
- [53] B. Meenakshi, A. Bhatnagar, and S. Roy. Tool for Translating Simulink Models into Input Language of a Model Checker. In *Formal Methods and Software Engineering*, volume 4260 of *LNCS*, pages 606–620. Springer, 2006.
- [54] B. Meyer. Applying “Design by Contract”. *Computer*, 25(10):40–51, 1992.
- [55] S. Minopoli and G. Frehse. SL2SX Translator: From Simulink to SpaceEx Verification Tool. In *19th ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, 2016.
- [56] K. S. Namjoshi and R. J. Trefler. On the completeness of compositional reasoning methods. *ACM Trans. Comput. Logic*, 11(3), 2010.
- [57] O. Nierstrasz. Regular types for active objects. *SIGPLAN Not.*, 28(10):1–15, 1993.
- [58] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*. LNCS 2283. Springer, 2002.
- [59] A. Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008.
- [60] A. Pnueli. The Temporal Logic of Programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977.
- [61] V. Preoteasa, I. Dragomir, and S. Tripakis. A Nondeterministic and Abstract Algorithm for Translating Hierarchical Block Diagrams. *CoRR*, abs/1611.01337, Nov. 2016.
- [62] V. Preoteasa, I. Dragomir, and S. Tripakis. A nondeterministic and abstract algorithm for translating hierarchical block diagrams. *CoRR*, abs/1611.01337, 2016.
- [63] V. Preoteasa, I. Dragomir, and S. Tripakis. Type Inference of Simulink Hierarchical Block Diagrams in Isabelle. In *37th IFIP WG 6.1 International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE)*, 2017.
- [64] V. Preoteasa and S. Tripakis. Refinement calculus of reactive systems. In *Embedded Software (EMSOFT), 2014 International Conference on*, pages 1–10, Oct 2014.

- [65] V. Preoteasa and S. Tripakis. Towards Compositional Feedback in Non-Deterministic and Non-Input-Receptive Systems. In *31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2016.
- [66] G. Reissig, A. Weber, and M. Rungger. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Trans. Automat. Contr.*, 62(4):1781–1796, 2017.
- [67] P. M. Rondon, M. Kawaguci, and R. Jhala. Liquid types. In *29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI ’08, pages 159–169, New York, NY, USA, 2008. ACM.
- [68] P. Roy and N. Shankar. SimCheck: An expressive type system for Simulink. In C. Muñoz, editor, *2nd NASA Formal Methods Symposium (NFM 2010), NASA/CP-2010-216215*, pages 149–160, Langley Research Center, Hampton VA 23681-2199, USA, Apr. 2010. NASA.
- [69] S. Sankaranarayanan and A. Tiwari. Relational abstractions for continuous and hybrid systems. In *Computer Aided Verification: 23rd International Conference, CAV 2011*, pages 686–702. Springer, 2011.
- [70] H. Schmeck. Algebraic characterization of reducible flowcharts. *Journal of Computer and System Sciences*, 27(2):165 – 199, 1983.
- [71] V. Sfyrla, G. Tsiligiannis, I. Safaka, M. Bozga, and J. Sifakis. Compositional translation of Simulink models into synchronous BIP. In *Industrial Embedded Systems (SIES), 2010 International Symposium on*, pages 217–220, July 2010.
- [72] N. Shankar. Lazy compositional verification. In *COMPOS’97: Revised Lectures from the International Symposium on Compositionality: The Significant Difference*, pages 541–564, London, UK, 1998. Springer-Verlag.
- [73] A. Siirtola, S. Tripakis, and K. Heljanko. When do we not need complex assume-guarantee rules? *ACM Trans. Embed. Comput. Syst.*, 16(2):48:1–48:25, Jan. 2017.
- [74] A. P. Sistla, M. Y. Vardi, and P. Wolper. The Complementation Problem for Büchi Automata with Applications to Temporal Logic. *Theoretical Computer Science*, 49:217–237, 1987.
- [75] G. Stefanescu. *Network Algebra*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2000.
- [76] S. Tripakis, B. Lickly, T. Henzinger, and E. Lee. On Relational Interfaces. In *Proceedings of the 9th ACM & IEEE International Conference on Embedded Software (EMSOFT’09)*, pages 67–76. ACM, 2009.
- [77] S. Tripakis, B. Lickly, T. A. Henzinger, and E. A. Lee. A Theory of Synchronous Relational Interfaces. *ACM Trans. Program. Lang. Syst.*, 33(4):14:1–14:41, July 2011.
- [78] S. Tripakis, C. Sofronis, P. Caspi, and A. Curic. Translating Discrete-time Simulink to Lustre. *ACM Trans. Embed. Comput. Syst.*, 4(4):779–818, Nov. 2005.
- [79] N. Wirth. Program development by stepwise refinement. *Comm. ACM*, 14(4):221–227, 1971.
- [80] D. Yadav and M. Butler. Verification of liveness properties in distributed systems. In *Contemporary Computing*, volume 40 of *Communications in Computer and Information Science*, pages 625–636. Springer Berlin Heidelberg, 2009.
- [81] C. Yang and V. Vyatkin. Transformation of Simulink models to IEC 61499 Function Blocks for verification of distributed control systems. *Control Engineering Practice*, 20(12):1259–1269, 2012.
- [82] C. Zhou and R. Kumar. Semantic Translation of Simulink Diagrams to Input/Output Extended Finite Automata. *Discrete Event Dynamic Systems*, 22(2):223–247, 2012.
- [83] L. Zou, N. Zhany, S. Wang, M. Franzle, and S. Qin. Verifying Simulink diagrams via a Hybrid Hoare Logic Prover. In *Embedded Software (EMSOFT)*, Sept 2013.