

Reactive Functional Programming *

R. B. Kieburtz

Oregon Graduate Institute of Science & Technology

P.O. Box 91000, Portland, OR 97291-1000 USA dick@cse.ogi.edu

Abstract

Reactive systems respond to concurrent, possibly unsynchronized streams of input events. Programming reactive systems is challenging without language support for event-triggered actions. It is even more challenging to reason about reactive systems. This paper explores a new conceptual basis for applying functional programming techniques to the design and formal verification of reactive systems. The mathematical foundation for this approach is based upon signature coalgebras and derived proof rules for coinduction. The concepts are illustrated with an example that has been used with the language *Esterel*.

Keywords

Reactive systems, programming languages, coalgebras, coinduction, verification

1 INTRODUCTION

Reactive systems are characterized by sequences of history-determined reactions to external events. It is known that a non-strict functional programming language can provide a suitable linguistic vehicle for programming reactive systems because streams, modeling temporal sequences of values, can be represented. It is necessary to represent more than streams, however. Current reactive programming languages, such as *Esterel*, *Lustre*, *Signal* and *Statecharts* provide implicit or explicit representations of state, iterative control structures, and parallel threads of activity. Use of these languages has advanced the state of the art of designing reactive systems, however, it is not easy to reason about their properties. For ease of reasoning, we should like to have a sound programming logic that is expressive over the terms of the programming language. A principal motivation for this research is to develop in tandem a programming notation well suited to specifying reactive systems, and an associated programming logic.

The control structure needed for reactive programs is inherently iterative, not recursive. The data of interest are infinitary sequences or trees of states, representing the evolution of systems that may never terminate. We have searched for an underlying mathematical structure to model reactive systems. The structure we have found most useful is that of coalgebras, which unfortunately, are not very well understood by most functional programmers.

*The research reported in this paper was supported by the USAF Materiel Command.

Section 2 of the paper introduces the notation and concepts of programming with coalgebras. Section 3 illustrates application of the concepts to formulate a non-trivial example that has previously appeared in the literature. This example is used to illustrate one of the main points of the paper: that a coalgebra provides structure for verification of properties as well for an algorithm. Section 4 presents brief conclusions.

1.1 Iterative functional programming

In the past, several researchers have observed that programming with terminating recursion schemes is algebraic in nature [7, 8, 4, 13, 23, 22]. Results of this body of research include (1) the discovery that recursion over typed data structures has a logical counterpart in structural induction and (2) that the recursion patterns can be captured in higher-order, type-parametric (also called *polytypic*) combinators, embedded in a strict functional programming language. There is a dual to algebraic programming and it leads to another style of functional programming, which is the topic of this paper.

Process-oriented programs are iterative. They are controlled by tests of their partial results, driven by external events rather than by the interpretation of data, and are naturally modeled by coalgebras. Because control is derived to meet external demand rather than induced by the structure of arguments, a non-strict evaluation mechanism is needed. The rules of coinduction induced by codatatypes are logical duals of the familiar rules of structural induction that are induced by datatype definitions. Dual to the recursive structure of algebraic programs is the iterative structure of coalgebraic programs, which can be made manifest by embedding a set of type-parametric combinators in a non-strict programming language.

Examples of iterative algorithms are common. They include linear and tree-structured searching, shift-reduce parsing, and both synchronous and asynchronous reactive systems. Before looking at examples, including proof rules, let's introduce a formalism for expressing coalgebraic programs. The notation presented in this paper is taken from *DUALITY*, which is an experimental, functional language based upon algebras and coalgebras as its fundamental computational structures. In this paper we shall deal only with the coalgebraic part. An early version of this language has been implemented and is described in a technical report [21].

This is by no means the first time that the concepts of iterative control and synchronous data flow (i.e. streams) have been embedded into a programming notation. These concepts were the basis for the language *Lucid* [2], and earlier, in Kahn-McQueen networks [20]. These early language developments led directly to *Lustre* [15], a language that has found considerable use for programming reactive systems. Synchronous data flow techniques have been further generalized [9] to networks of higher-order typed entities.

2 COVARIETIES OF COALGEBRAS

A covariety is a class of coalgebras with a common signature. The archetypical example is the covariety of stream coalgebras, whose signature is:

cosig $Stream(a)\{\mathbf{type} \ c; \ str/c : \{\$shd : a, \$stl : c\}\};$

Here $Stream$ is the name of the covariety and str is the name of its single sort, analogous to a type constructor. It is redundant to name both the covariety and its single sort, but later we shall see an example that requires a multi-sorted signature.

Each $Stream$ -coalgebra has a type parameter, a , a carrier type, c , and two projectors (some authors call these *destructors*) identified by the symbols $\$shd$ and $\$stl$. Each projector is a total function whose domain is the carrier and whose codomain is indicated by the typing given in the signature. A signature of projectors can be thought of as a generalized record declaration. Binding a type for the carrier and typed functions for the projectors defines a specific coalgebra parameterized by the type variable, a .

Every covariety defined in this way contains a final coalgebra which is unique up to isomorphism [18]. Call the carrier of the final $Stream$ -coalgebra $str(a)$ and call its projectors Shd and Stl . The finality condition implies that given any coalgebra of the variety $Stream$, with bindings $\{c := t, \$shd := f, \$stl := g\}$, there is an assignment of the type parameter, $a := t'$, and a unique mapping, $h : t \rightarrow str(t')$, satisfying the homomorphism condition:

$$Stl \circ h = h \circ g$$

$$Shd \circ h = f$$

The significance of the final coalgebra is that any $Stream$ -coalgebra may be represented as an infinite sequence of values. Because of this property, a stream is often identified with an infinite list. While that is certainly one way to encode a stream as data, it is by no means the only way. If we accept that a stream is *codata*, then encoding it as data seems unnecessary. A codata object is defined with methods for observing it, rather than with constructors for building its representation. Other examples of codata are found in the *intensions* of *Lucid*, which denote suspended, iterative computations, rather than values.

Every stream is infinite; that is, it is meaningful to iterate the projection operator Stl on a stream arbitrarily many times, even though there is no way to witness the entire stream at once. A stream provides a good model for an incrementally readable input file. The projection Shd yields the value of the first element of a stream, just as a *get* operation on an open file produces a value from it. The projection Stl yields the rest of a stream, but it is not

manifested until projections of it are taken. The situation is familiar in non-strict functional languages.

There are infinitely many access paths to elements of a stream. A path is expressed by a well-typed composition of the projectors Shd and Stl , i.e. $Shd \circ \underbrace{Stl \circ \dots \circ Stl}_{i \text{ times}}$, for $i \geq 0$.

2.1 Generators of codata

We say that the carrier of a final coalgebra is a type of codata, meaning that data values can be gotten from it by projection. Data and codata are distinguished by the type system of *DUALITY*. When a sort of a covariety is used as a type constructor, it designates a type of codata.

If $T(a)$ is a covariety and t is a sort symbol of the signature T , a generator of sort t is a function with a type $c \rightarrow t(a)$, where the codomain is a type of codata. To determine a generator, we must specify a coalgebra by naming a covariety, a type for the carrier and bindings of functions for the projectors. Coalgebras are first-class objects of *DUALITY*.

Example 21 : Integer sequences

Declare a *Stream* coalgebra by:

```
coalgebra intseq  $\doteq$  Stream{c := int; $shd := id, $stl := add1}
```

where $add1 \doteq \lambda n. n + 1$. To obtain an expression of type $str(int)$, we can apply the *DUALITY* combinator, $gen[str]$, to the coalgebra specification, creating a generator for this type,

```
gen[str] intseq : int  $\rightarrow$  str(int)
```

The combinator $gen[str]$ is an instance of a higher-order combinator, gen , specialized to the sort str of the covariety *Stream*. The higher-order combinators in *DUALITY* substitute for the recursion operator found in conventional functional languages. The generator $h \doteq gen[str] intseq$ is the unique map taking the coalgebra $intseq$ to a final *Stream* coalgebra. The homomorphism condition it satisfies is expressed by the pair of equations:

$$Shd \circ h = id_{int}$$

$$Stl \circ h = h \circ add1$$

□

Other applications of *Stream* coalgebra generators define pseudo-random number sequences, sequences of unique identifiers and other enumerated sets.

(a) A proof rule for stream generators

All functional programmers are familiar with proof rules based upon induction. Somewhat less familiar is the dual rule of coinduction. The possible observations of a codata object are enumerable composites of a finite basis of primitive witness functions. The coinduction principle is that the finitely observable properties of an object completely characterize it, even if the object is infinitary [1, 24, 26, 18, 12]. Expositions of coinduction are given by Paulson [25], Gordon [14], Jacobs and Rutten [19].

To define a proof rule for a stream of elements of type a , generated from a carrier of type t , let P be a two-place, typed predicate symbol whose arguments range over t and a , respectively. We prefer a two-place predicate because it can express the input-output relation of a function. Coinduction extends the domain of the relation to infinitary objects. A proof rule for a stream of elements is:

$$\frac{\begin{array}{l} x_0 : t \\ f : t \rightarrow a \\ g : t \rightarrow t. \quad \forall x : t. P(x_0, x) \Rightarrow P(x_0, g x) \end{array}}{P(x_0, x_0) \Rightarrow \Box P(x_0, \text{gen}[\text{str}]\{c := t; \$shd := f, \$stl := g\} x_0)}$$

We have used a linear temporal operator, \Box (read as *always*), as a quantifier on the predicate P in the consequent of the rule to express that the proposition $P(x_0, x)$ is asserted for every element, x , of the generated stream.

2.2 Coalgebra homomorphisms define iteration schemes

A compelling reason to consider coalgebras is that *coalgebra homomorphisms*, i.e. the structure-preserving maps between coalgebras of a given covariety, conform directly to an iteration scheme for computation. Thus coalgebras afford a mechanism to prescribe a specific control structure for computation and to communicate this structure to program analysis and translation software.

For the covariety *Stream*, the related iteration scheme is linear search. For more complex covarieties, the iteration schemes are more specialized, including algorithm schemes such as binary search and shift-reduce parsing.

A coalgebra homomorphism is composed of two parts: a coalgebra specification, such as *intseq* in Example 21, and a *control* that selects among the projectors of the coalgebra. The body of a control has the form of a conditional or a case expression. A control and a coalgebra specification are combined by a *DUALITY* combinator, *cohom*, suitably specialized to a sort of the covariety.

This forms a limit of the specified coalgebra, determined by the control. It is, of course, necessary to confirm that such limits exist, in each case.

Example 22 : Sequential search

We shall define a generic sequential search function, give a necessary and sufficient condition for its termination, and give a hypothetical rule of logic to conclude a property of a search. Let a be a type, $p : a \rightarrow \text{bool}$ and $r : a \rightarrow a$. Define a sequential search combinator, *while*, by

$$\begin{aligned} \text{while}(p, r) \doteq \text{cohom}[\text{str}] \text{Stream} \{c := a; \$shd := id_a, \$stl := r\} \\ (\lambda x. \text{let } u = \$shd \ x \ \text{in} \\ \quad \text{if } \neg p \ u \ \text{then } u \ \text{else } \$stl \ x) \end{aligned}$$

The control is expressed as a lambda abstraction enclosed in parentheses. The expressions on the arms of a conditional (or case expression) that forms the return expression of the control must be applications of projectors of the coalgebraic variety, or as in this example, identifiers bound to such applications in a local definition. A control expression should not be confused with a function declaration; in particular, the types of the expressions on the arms of a conditional or case are not all of a common type.

A function composed with $\text{cohom}[\text{str}]$ satisfies a set of conditional equations such as the ones given below for the sequential search combinator:

$$\begin{aligned} p \ x = tt \quad \Rightarrow \quad \text{while}(p, r) \ x = \text{while}(p, r) \ (r \ x) \\ p \ x = ff \quad \Rightarrow \quad \text{while}(p, r) \ x = id_a \ x \end{aligned}$$

The right-hand sides of the equations are formed by substitution into the control expression. The bindings of the projectors $\$shd$ and $\$stl$ are taken from the coalgebra declaration, and in addition, the defined combinator is recursively applied to every projector expression whose codomain type has been specified to be the carrier. In this example, the declaration $\$stl : c$ in the signature dictates that $\text{while}(p, r)$ is applied to $(r \ x)$, gotten from the binding of $\$stl$. It is not misleading to imagine the combinator expressions of *DUALITY* translated in this way into recursive function definitions in a conventional language. However, the patterns of recursion so obtained are rigidly constrained to tail recursion.

Of course, it is useful to have a more compact declaration for such a useful combinator as $\text{while}(_, _)$ and it is often made a language primitive. We have used it here as the simplest illustration of coalgebraic program construction with explicit control. Before leaving the example, we call the reader's attention to another aspect of the coalgebraic declaration. The declarations of data transformation and control are separated; each is a first-class entity in *DUALITY*. The data transformation is fully specified by the coalgebra, which

could be used in other declarations with a different control. The control specification could be used with other *Stream* coalgebras.

3 FINITE-STATE REACTIVE SYSTEMS

Finite-state systems are naturally modeled by multi-sorted coalgebras. The states of a system correspond to sorts of a coalgebra; the carrier in each sort is comprised of the state variables, and the projectors in each sort are the possible transitions from the state that corresponds to the sort. Traditional functional programming languages have not been easy to use in describing reactive systems because the sequences of possible reactions often seem to require complex mutual recursion for their specification. Formulating a reactive system as a coalgebra is more natural because the use of multiple sorts provides a detailed structure for the specification.

We shall illustrate the technique with an example of a synchronous reactive system previously used to illustrate programming in *Esterel* [3].

Example 31 : The Reflex game

The Reflex game is a coin-operated machine on which a player measures the time constant of her reflexes. After depositing a coin to start the game, she can depress a Ready button to signify that she is prepared to start a trial. When she receives a Go signal from the machine, she depresses a Stop button as quickly as she can. The machine times her response in several trials, then displays the average response time. There are several illegal moves that must be accounted for. If the Stop button is depressed after the player is ready but before Go has been signaled, this action is interpreted as cheating and terminates the game. If either the Ready or the Stop button is depressed when it is not expected, a warning bell sounds, but the game is not interrupted. A coin drop always restarts the game, even when this event occurs during the progress of a previous game.

The game also depends upon timing signals emitted by a clock. Clock ticks must be counted to measure the player's latency. Also, the Go signal is emitted after a randomly determined number of clock ticks following depression of the Ready button by the player. And if a player fails to respond within a predetermined interval when a response is expected, the game times out.

The events that the machine must react to are a coin drop, depression of the Ready and Stop buttons, and ticks of the clock. We assume that these events never occur exactly simultaneously, or that they can be separated in a sequence.

Analysis of the Reflex game shows that the machine can be described as having five major states:

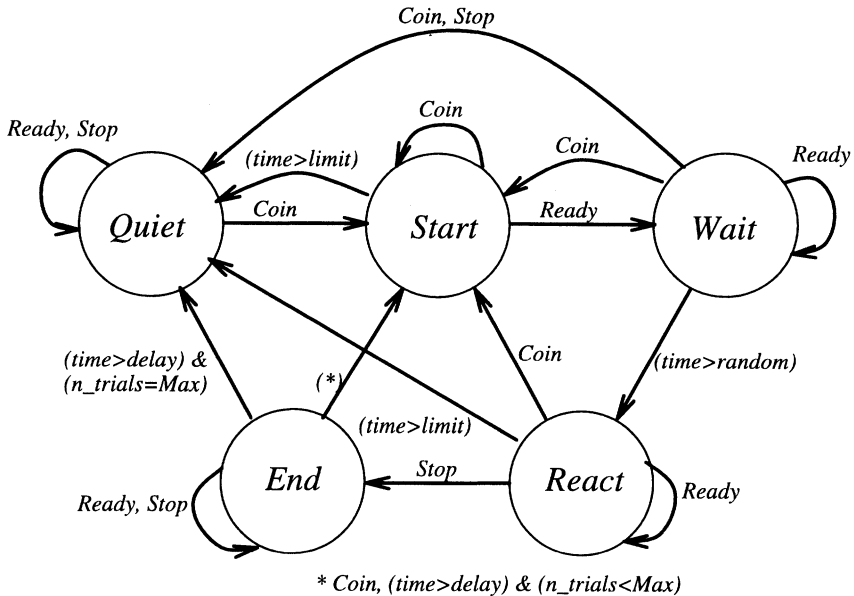


Figure 1—Major states of the Reflex game

- quiet*, when no game is in progress,
start, when awaiting a Ready event to start a trial,
wait, when the player is awaiting a Go signal from the machine,
react, when the machine awaits a Stop event,
end, when the machine pauses to display the response time of the player.

The machine responds to events differently in each of these five states. Some of the responses are transitions from one state to another. Figure 1 is a state transition diagram for the reflex game machine.

In the solution of this problem as an *Esterel* program, the states of the game are not manifest but are implicit in the control. The control consists of a nested loop structure, triggered by events, that takes the machine through the possible sequences of state transitions. Although *Esterel* provides intuitive syntax for coding even-driven nested loop structures, it is still challenging to get them right. This represents the state-of-the-art in programming reactive systems.

The Reflex game can be modeled by a multi-sorted coalgebra. We associate a separate sort with each of the major states of the game. In each of these states, we identify the possible transitions and name them. The reactions in each state become the projectors of the corresponding sort. The codomain type of each projector is the carrier that corresponds to the game state to which the transition leads. No explicit recursion or iteration is involved in programming the game in this way.

The output of the Reflex game will be modeled as a sequence of states. A state will include state variables and output signals produced by a transition. However, these are details that will appear in a coalgebra for the game. None of these details are manifested in a covariety. A signature for the covariety is:

```
cosig Reflex(a) {type q, s, w, r, t;
  quiet/q : {$coin : s, $noop : q},
  start/s : {$reveal : a, $ready : w, $renew : s, $warn : s, $timeout : q, $tick : s},
  wait/w : {$reveal : a, $renew : s, $warn : w, $abort : q, $tick : w, $go : r},
  react/r : {$reveal : a, $react : t, $renew : s, $warn : r, $tick : r, $timeout : q},
  end/t : {$reveal : a, $warn : t, $renew : s, $tick : t, $stock : s, $finish : q}}
```

The projectors *\$reveal* do not correspond to state transitions of the game, but are instead actual projections of the machine state.

The game is specified in terms of a *Reflex*-coalgebra, binding data transformation functions associated with state transitions to each of the projector symbols. To determine all trajectories of play, we shall generate a game tree, which will be codata, of course. To simulate a game, we shall interpret a sequence of externally caused events (coin drops, clock ticks and button pushes) as control for the projectors in each state.

Minor states of the game are determined by the values of three integer-valued state components. These are packaged as fields of a record type,

```
record State{time, total_time, trial_number : int};
```

Further, there are signals delivered to the actuators that implement the machine. These unvalued signals can be represented by a set of elements of an enumerated type,

```
type Signals = set of [game_over_on | game_over_off | go_on | go_off |
  tilt_on | tilt_off | ring_bell | bump_random];
```

There is one integer-valued signal which sends values to the display. The state components and the signals are conveniently packaged as fields of a record type. We declare

```
record Game{state : State; sigs : Signals; display : int};
```

It is convenient to define a pair of constants of type *Game*,

```
def initial_game = {state := {time := 0; total_time := 0; trial_number := 0};
  sigs := [game_over_off, go_off, tilt_off];
  display := 0};
def tilt_game = {{state := {time := 0; total_time := 0; trial_number := 0};
  sigs := [game_over_on, go_off, tilt_on];
  display := 0};
```

There are also three integer constants, *Time_limit*, *Delay* and *Max_trials*, and a stream, *random*, which is a randomly generated sequence of positive integers of bounded size, supplied by the machine. As a notational abbreviation, let $R \oplus \{X := e\}$ denote the record whose fields have the values of the corresponding fields in the record R , except for field X , which has the value of e .

The next task is to define a coalgebra by specifying the projectors of each sort. These correspond to the possible transitions from each state.

```

reflex  $\doteq$  coalgebra Reflex { $q, s, w, r, t :=$  state,
  quiet : { $\$coin := \lambda s.$  initial_game,
            $\$noop := id_{Game}$ },
  start : { $\$reveal := \lambda s.$  s.state,
            $\$ready := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := 0},
                                sigs := [bump_random]}},
            $\$renew := \lambda s.$  initial_game,
            $\$warn := \lambda s.$  s  $\oplus$  {sigs := [ring_bell]},
            $\$timeout := \lambda s.$  tilt_game,
            $\$tick := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := s.state.time + 1}}},
  wait : { $\$reveal := \lambda s.$  s.state,
            $\$renew := \lambda s.$  initial_game,
            $\$warn := \lambda s.$  s  $\oplus$  {sigs := [ring_bell]},
            $\$abort := \lambda s.$  tilt_game,
            $\$tick := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := s.state.time + 1}},
            $\$go := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := 0}; sigs := [go_on]}},

  react : { $\$reveal := \lambda s.$  s.state,
            $\$react := \lambda s.$  {state := {time := 0;
                                total_time := s.state.total_time + s.state.time;
                                trial_number := s.state.trial_number + 1};
                                sigs := [go_off]; display := s.state.time},
            $\$renew := \lambda s.$  initial_game,
            $\$warn := \lambda s.$  s  $\oplus$  {sigs := [ring_bell]},
            $\$tick := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := s.state.time + 1}},
            $\$timeout := \lambda s.$  tilt_game},
  end : { $\$reveal := \lambda s.$  s.state,
          $\$renew := \lambda s.$  initial_game,
          $\$warn := \lambda s.$  s  $\oplus$  {sigs := [ring_bell]},
          $\$tick := \lambda s.$  s  $\oplus$  {state  $\oplus$  {time := s.state.time + 1}},
          $\$stock := \lambda s.$  {state := {time := 0},
                           display := 0},
          $\$finish := \lambda s.$  s  $\oplus$  {sigs := [game_over_on];
                                display := s.state.total_time/Max_trials}}}
```

A generator composed from this coalgebra, for instance, $gen[quiet] reflex : Game \rightarrow quiet$ generates, in response to demand, an infinite game tree rooted on the quiescent game state. Paths in this game tree incorporate all major-

state transitions allowed by the rules of the game and in addition, some that are not allowed, because transitions in the game tree are unconstrained by conditions on the state variables that govern the progress of an actual game (i.e. the rules of the game).

To obtain a function that accurately simulates the game, the coalgebra must be composed with a control that responds to input events and reads state variables to determine a game path. The control is defined as a cluster of five expressions, one for each sort, as the game's response to events depends upon the major state that it occupies.

Since the final result is a stream of game states, the body of each component of the control has the form of a *Stream*-generator. The *\$shd* projector defined in each game state translates each of the possible input events into a state transition event. In some cases, the translation is conditioned by the elapsed time recorded in a game state. Here is a definition of the control:

```
def transition =
  (quiet : (λs. gen[str] Stream{c := str(event);
    $shd := λes. let e = Shd es in
      case e of
        Coin ⇒ $coins
        | Ready ⇒ $noops
        | Stop ⇒ $noops
        | Tick ⇒ $noops
      end
    $stl := Stl}),
  start : (λs. gen[str] Stream{c := str(event);
    $shd := λes. let e = Shd es in
      case e of
        Coin ⇒ $renew s
        | Ready ⇒ $ready s
        | Stop ⇒ $warn s
        | Tick ⇒ let v = $reveals in
          if v.time < Time_limit then $tick s
          else $timeout s
        end
      $stl := Stl}),
  wait : (λs. gen[str] Stream{c := str(event);
    $shd := λes. let e = Shd es in
      case e of
        Coin ⇒ $renew s
        | Ready ⇒ $warn s
        | Stop ⇒ $abort s
        | Tick ⇒ let v = $reveals in
          if v.time < random then $tick s
          else $go s
        end
      $stl := Stl}),
```

```

react : (λs. gen[str] Stream{c := str(event);
    $shd := λes. let e = Shd es in
        case e of
            Coin ⇒ $renew s
          | Ready ⇒ $warn s
          | Stop ⇒ $react s
          | Tick ⇒ let v = $reveals in
                    if v.time < Time.limit
                    then $tick s else $timeout s
        end
    $stl := Stl}),
end : (λs. gen[str] Stream{c := str(event);
    $shd := λes. let e = Shd es in
        case e of
            Coin ⇒ $renew s
          | Ready ⇒ $warn s
          | Stop ⇒ $warn s
          | Tick ⇒ let v = $reveals in
                    if v.time < Delay then $tick s
                    else let v = $reveals in
                           if v.trial_number < Max.trials
                           then $tock s
                           else $finish s
                    end
        end
    $stl := Stl}))

```

A simulator for the Reflex game is the function

cohom[quiet] reflex transition : Game → str(event) → str(Game).

□

The domain constraints needed for the example of the Reflex game can be established by structural (Hindley-Milner) type checking. Note that the simulator is not expected to terminate, in the usual sense, but rather to make finite progress in response to each external event that it receives.

Programming the Reflex game in terms of coalgebras is straightforward once the type of the solution and of the component functions has been determined. The structure of the coalgebra does not allow guesswork.

3.1 A verification logic for the Reflex game simulator

One of the most significant advantages of formulating a finite-state systems such as the Reflex game simulator as a coalgebra morphism is that the coalgebraic structure induces a complementary deductive logic in which properties of the system can be proved by coinduction. As we shall see, the coinduction

rules induced by the coalgebraic structure provide a fine-grained decomposition of proof obligations that must be discharged to establish a conjectured property. We believe this structure will make verification significantly easier by removing most of the guesswork. We expect it to be amenable to the application of automatic proof discovery methods.

Corresponding to each carrier in the coalgebra signature declaration, we declare a predicate symbol whose interpretation will characterize a specific property in the major state (or sort) to which the carrier is bound. For the reflex game example, these will be unary* predicates, each relating an external event stream and a game state. The game state will be a minor state of the major state characterized by the predicate.

A coinduction rule for a coalgebra is formulated as a sequent clause. In the consequent are clauses for each sort of a multi-sorted coalgebra; in the antecedent are sets of hypotheses for each sort. The hypotheses for a given sort correspond one-for-one to the projectors defined for that sort. The structure of a coinduction rule is induced directly by the signature of a coalgebra.

Each clause in the consequent of a coinduction rule extends the interpretation of a predicate to encompass all of the states of the corresponding sort in a potentially infinite tree or sequence. There will be one such clause for each sort of a multi-sorted coalgebra, allowing characterization of a property specified at each of the major states of a finite-state model, throughout all minor states that are reachable from a given initial state.

Each individual clause of an antecedent implies the transfer of a property under a projection. For instance, referring to the Reflex game, a clause that implies the transfer of a property via the transition *\$ready* is:

$$\forall u : s, es : str(event). S(u) \Rightarrow W(\$ready\ u)$$

where S and W are the predicate symbols associated with sorts s and w , respectively. Upon substituting the binding for the transition *\$ready* as given in the declaration of the coalgebra *reflex*, the clause becomes

$$\forall u : s, es : str(event). S(u) \Rightarrow W(u \oplus \{state \oplus \{time := 0\}\})$$

When the coinduction rule is for a general coalgebra morphism (a *cohom*), each hypothetical implication must be qualified by a guard for the transition that can be read from the declaration of the control for the morphism. Again referring to the reflex game, the clause above, extended as a hypothetical clause for the simulator, becomes

*In this example, the initial state is fixed in each game. More generally, use of a two-place predicate would allow the simulator to be characterised as a function from an arbitrarily specified initial state to the ensuing behavior.

$$\begin{aligned}
& \forall u : q, es : \text{str}(\text{event}). (\text{Shd } es = \text{Coin}) \Rightarrow Q(u) \Rightarrow S(\text{initial_state}) \\
& \forall u : q, es : \text{str}(\text{event}). (\text{Shd } es = \text{Ready}) \Rightarrow Q(u) \Rightarrow Q(u) \\
& \forall u : q, es : \text{str}(\text{event}). (\text{Shd } es = \text{Stop}) \Rightarrow Q(u) \Rightarrow Q(u) \\
& \forall u : q, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow Q(u) \Rightarrow Q(u) \\
& \forall u : s, es : \text{str}(\text{event}). (\text{Shd } es = \text{Coin}) \Rightarrow S(u) \Rightarrow S(\text{initial_state}) \\
& \forall u : s, es : \text{str}(\text{event}). (\text{Shd } es = \text{Ready}) \Rightarrow S(es, u) \Rightarrow W(u \oplus \{\text{state} \oplus \{\text{time} := 0\}\}) \\
& \forall u : s, es : \text{str}(\text{event}). (\text{Shd } es = \text{Stop}) \Rightarrow S(es, u) \Rightarrow S(u \oplus \{\text{sigs} := [\text{ring_bell}]\}) \\
& \forall u : s, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} < \text{Time.limit}) \Rightarrow \\
& \quad S(u) \Rightarrow S(u \oplus \{\text{state} \oplus \{\text{time} := s.\text{time} + 1\}\}) \\
& \forall u : s, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} \geq \text{Time.limit}) \Rightarrow \\
& \quad S(u) \Rightarrow Q(\text{tilt_game}) \\
& \forall u : w, es : \text{str}(\text{event}). (\text{Shd } es = \text{Coin}) \Rightarrow W(u) \Rightarrow S(\text{initial_state}) \\
& \forall u : w, es : \text{str}(\text{event}). (\text{Shd } es = \text{Ready}) \Rightarrow W(u) \Rightarrow W(u \oplus \{\text{sigs} := [\text{ring_bell}]\}) \\
& \forall u : w, es : \text{str}(\text{event}). (\text{Shd } es = \text{Stop}) \Rightarrow W(u) \Rightarrow Q(\text{tilt_game}) \\
& \forall u : w, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} < C.\text{random}) \Rightarrow \\
& \quad W(u) \Rightarrow W(u \oplus \{\text{state} \oplus \{\text{time} := s.\text{time} + 1\}\}) \\
& \forall u : w, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} \geq C.\text{random}) \Rightarrow \\
& \quad W(u) \Rightarrow W(u \oplus \{\text{state} \oplus \{\text{time} := 0\}; \text{sigs} := [\text{go_on}]\}) \\
& \forall u : r, es : \text{str}(\text{event}). (\text{Shd } es = \text{Coin}) \Rightarrow R(u) \Rightarrow S(\text{initial_state}) \\
& \forall u : r, es : \text{str}(\text{event}). (\text{Shd } es = \text{Ready}) \Rightarrow R(u) \Rightarrow R(u \oplus \{\text{sigs} := [\text{go_on}]\}) \\
& \forall u : r, es : \text{str}(\text{event}). (\text{Shd } es = \text{Stop}) \Rightarrow R(u) \Rightarrow \\
& \quad S(u \oplus \{\text{state} \oplus \{\text{time} := 0; \\
& \quad \quad \quad \text{total.time} := s.\text{state}.\text{total.time} + s.\text{state}.\text{time}; \\
& \quad \quad \quad \text{trial.number} := s.\text{state}.\text{trial.number} + 1\}; \\
& \quad \quad \quad \text{sigs} := [\text{go_off}]; \text{display} := s.\text{state}.\text{time}\}) \\
& \forall u : r, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} < \text{Time.limit}) \Rightarrow \\
& \quad R(u) \Rightarrow R(u \oplus \{\text{state} \oplus \{\text{time} := u.\text{time} + 1\}\}) \\
& \forall u : r, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} \geq \text{Time.limit}) \Rightarrow \\
& \quad R(u) \Rightarrow Q(\text{tilt_game}) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Coin}) \Rightarrow T(u) \Rightarrow S(\text{initial_state}) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Ready}) \Rightarrow T(u) \Rightarrow T(u \oplus \{\text{sigs} := [\text{ring_bell}]\}) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Stop}) \Rightarrow T(u) \Rightarrow T(u \oplus \{\text{sigs} := [\text{ring_bell}]\}) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} < \text{Delay}) \Rightarrow \\
& \quad T(u) \Rightarrow T(u \oplus \{\text{state} \oplus \{\text{time} := s.\text{time} + 1\}\}) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} \geq \text{Delay}) \wedge \\
& \quad (u.\text{trial.number} < \text{Max.trials}) \Rightarrow \\
& \quad T(u) \Rightarrow S(u \oplus \{\text{state} \oplus \{\text{time} := 0\} \\
& \quad \quad \quad \text{display} := 0\})) \\
& \forall u : t, es : \text{str}(\text{event}). (\text{Shd } es = \text{Tick}) \Rightarrow (u.\text{state}.\text{time} \geq \text{Delay}) \wedge \\
& \quad (u.\text{trial.number} \geq \text{Max.trials}) \Rightarrow \\
& \quad T(u) \Rightarrow Q(u \oplus \{\text{sigs} := [\text{game_over_on}]; \\
& \quad \quad \quad \text{display} := s.\text{state}.\text{total.time}/\text{Max.trials}\}))
\end{aligned}$$

Figure 2—Antecedent clauses for the Reflex coinduction rule

$$\forall u : s, es : str(event). (transition.start\ u\ (Shd\ es) = \$ready\ u) \Rightarrow S(u) \Rightarrow W(u \oplus \{state \oplus \{time := 0\}\})$$

Further substituting the guard clause by its binding in the declaration of *transition*, the hypothetical implication now relates the transition to the occurrence of an external event:

$$\forall u : s, es : str(event). (Shd\ es = Ready) \Rightarrow S(u) \Rightarrow W(u \oplus \{state \oplus \{time := 0\}\})$$

This recipe gives us the following coinduction rule for the Reflex game simulator. The antecedent clauses are given in Figure 2. The consequents are:

$$\begin{aligned} \forall u_0 : Game, es : str[event]. Q(u_0) &\Rightarrow \Box Q(cohom[quiet]\ reflex\ transition\ u_0\ es) \\ \forall u_0 : Game, es : str[event]. S(u_0) &\Rightarrow \Box S(cohom[start]\ reflex\ transition\ u_0\ es) \\ \forall u_0 : Game, es : str[event]. W(u_0) &\Rightarrow \Box W(cohom[wait]\ reflex\ transition\ u_0\ es) \\ \forall u_0 : Game, es : str[event]. R(u_0) &\Rightarrow \Box R(cohom[react]\ reflex\ transition\ u_0\ es) \\ \forall u_0 : Game, es : str[event]. T(u_0) &\Rightarrow \Box T(cohom[end]\ reflex\ transition\ u_0\ es) \end{aligned}$$

The temporal operator \Box (*always*) in the consequent formulas expresses precisely the sense in which the predicate over game states is extended by coinduction to a predicate over a stream of game states.

The textual extent of this coinduction rule is imposing, but keep in mind that it is amenable to mechanical calculation from the three parts of the formal declaration of the Reflex game simulator: the signature, the coalgebra specification and the control. When given a conjectured proposition of a property of the potentially infinite behaviors of the simulator, the rule yields a finite set of finitary propositions that must be discharged to prove the proposition. It accomplishes a logical destructuring that is essential to constructing a formal proof, and it does so in a way that suggests mechanization.

(a) Proving safety properties of the Reflex game

The consequents of the coinduction rule for the Reflex game assert invariant properties of states of the game. These are its so-called *safety* properties.

A simple safety property is that in every state, u ,

$$Q_t(u) \doteq \frac{C_random \leq Time_limit \quad Delay \leq Time_limit}{u.state.time \leq Time_limit}$$

In this clause, the antecedent conditions relate the values of constants of the Reflex game. Without these relations, the property does not hold. To prove this property of a game started in the quiescent state, we formulate the conjectured property as a consequent:

$$\forall u_0 : Game, es : str[event]. (u_0.state.time \leq Time_limit) \Rightarrow \Box (state.time \leq Time_limit)(cohom[quiet]\ reflex\ transition\ u_0\ es)$$

for which we seek a proof by Reflex game coinduction.

In the coinduction rule stated in the preceding section, choose $Q \equiv S \equiv W \equiv R \equiv T \equiv Q_t$ and attempt to discharge each of the antecedents of the rule. Most of the antecedent clauses discharge trivially, either because they do not refer to the *time* parameter explicitly or they set it to zero. There are five antecedent clauses in which the *time* parameter is incremented, however. Each of these clauses is guarded by a condition that *time* is strictly less than one of the constants *Time_limit*, *C_random* or *Delay*. Using the antecedent condition relating the latter two constants to *Time_limit*, it can be established that the implicand in each of the clauses is satisfied and the clause can be discharged. Thus the property is proved to hold for every game state reachable from an initial state that satisfies the property, by coinduction.

A related safety property that can be established is

$$Q_{tt}(u) \equiv u.state.total_time \leq Max_trials * Time_limit$$

We can also state safety properties consequent to a restriction on the external event stream. For instance we might assert

$$\begin{aligned} \forall u_0 : Game, es : str(event). \\ \square(next\ event \neq Coin)es \Rightarrow Q(u_0) \wedge (u_0.state.time = 0) \Rightarrow \\ \square(state.time = 0)(cohom[quiet]\ reflex\ transition\ u_0\ es) \end{aligned}$$

for a game started from the quiescent state. To prove this assertion, we must infer from the temporal logic assertion, $\square(next\ event \neq Coin) es$, the proposition $\forall es : str(event). Shd\ es \neq Coin$. This derived proposition can then be used to restrict the domain of antecedent clauses in a proof using the Reflex game coinduction rule.

Verification logics have also been developed for proving safety properties of systems specified in *Lustre* [17, 16] and *Esterel* [5].

3.2 Liveness properties of the Reflex game

Liveness properties, which assert that a state or state sequence with the property is eventually reached in every unfolding of the game, require for their proof a specific measure of progress. A monotonically increasing count of any external event can provide a suitable measure, provided that the timing event occurs *almost everywhere* in the event stream. Occuring almost everywhere means that at every point in the stream, the next timing event will occur after at most a finite number of non-timing events. The almost everywhere restriction assures that the observation of advancing time is never obscured by an infinite stream of non-timing events.

In the Reflex game, intuition tells us that the *Tick* event is a reasonable

choice for the timing event. We assume that it occurs almost everywhere in every possible stream of external events. This assumption is essential; it cannot be proved from weaker assumptions.

A liveness property we should like to prove is that every game started by a coin drop eventually terminates. We adopt for our definition of termination that either (a) the game enters the quiescent state or (b) another coin is dropped. Note that we have no direct characterization of the quiescent state (or any other major state) in terms of the program's state variables. The *time* attribute of the game state is only locally monotonic. It is monotonic with respect to transitions from any state to itself, but is reset to zero on many transitions from one state to another. Although the attribute *trial_number* is monotonic throughout all transitions that do not enter the quiescent state, we cannot use this attribute to characterize the quiescent state, because it is not incremented immediately upon leaving the quiescent state.

States can be characterized by sets of initial sequences of event streams, but this is not very convenient. It is more convenient to introduce a pseudo-variable, *game_over*:*bool*, which is set to *true* when the event *game_over_on* is signalled, and to *false* when *game_over_off* is signalled. The quiescent state is then characterized by a *true* value of *game_over*.

The liveness property asserting termination can be formally stated as:

$$\forall u_0 : \text{Game}, es : \text{str}(\text{event}). \Box(\text{next event} \neq \text{Coin})es \Rightarrow \\ \Diamond(\text{state.game_over} = \text{true})(\text{cohom}[\text{quiet}] \text{ reflex transition } u_0 \text{ es})$$

There is an accompanying safety property asserting that once a game terminates, *game_over* remains false if no more coins are dropped. We shall address only the liveness property.

The most successful way yet developed to verify temporal properties of a finite state system uses model checking of temporal logic formulas[11, 10]. The safety and liveness properties of the Reflex game example can obviously be verified by symbolic model checking. We describe a variant of the standard technique that uses symbolic inference to check monotonicity properties of state variables over transition paths. We have not yet implemented this method.

An ordered, symbolic binary decision diagram (BDD)[6] can be used to construct a proof of a liveness property. Boolean pseudo-variables are introduced to represent the boolean-typed expressions on which local control decisions are based. The nodes of the BDD correspond to major states of the Reflex game, split in cases in which more than one boolean condition controls transitions from the state. The boolean control expressions for each state are identified by inspecting the control specification.

(start) $x_1 \doteq state.time < Time.limit$
 (wait) $x_2 \doteq state.time < C.random$
 (react) $x_3 \doteq state.time < Time.limit$
 (end) $x_4 \doteq state.time < Delay$
 (end') $x_5 \doteq state.trial.number < Max.trials$

The pseudo-variables are ordered by $x_1 < x_2 < x_3 < x_4 < x_5$.

A nondeterministic BDD for the Reflex game is shown in Figure 3(a). Its nodes correspond to states of the Reflex game, but are labelled in the diagram by single letters S , W , R , E , E' to save space. The solid arcs indicate transitions possible when the value of the controlling pseudo-variable is positive; the dashed arcs represent transitions possible on negative values of the control variables.


Notice that there are multiple positive (or negative) arcs from some nodes. This BDD represents a nondeterministic FSA because transitions of the Reflex machine also depend upon external events which have not been represented in the control expressions bound to pseudo-variables. Note also that in constructing Figure 3(a), transitions that require the *Coin* event have been omitted because the *Coin* event is precluded by the antecedent clause of the liveness assertion. Nondeterminism allows us to represent with the BDD all of the transitions possible with event sequences that are restricted only by the assumptions that *Tick* events occur almost everywhere and *Coin* events are never present. The liveness property that we seek will be proved if we can show that the BDD of Figure 3(a) can be reduced to the single node, 1.

(a) Reducing a BDD with repeated nodes

We shall describe (informally) how to reduce the BDD of Figure 3(a), which represents the asserted liveness property of the Reflex game. Notice that this BDD contains paths from ancestor nodes to leaf nodes that carry the same labels. We call these *repeated nodes*. A path from an ancestor to a repeated node occurrence represents a loop in the state transition diagram. To reduce the BDD, each of these paths must be shown to be only finitely extensible as it is elaborated by repeating transitions of the state machine.

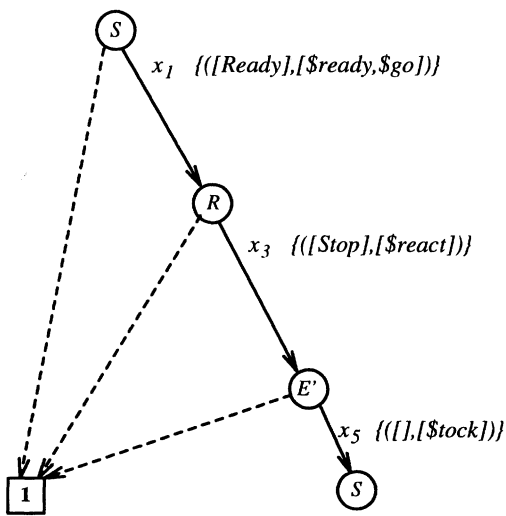
To establish that a path to a repeated node is only finitely extensible, we examine the control expressions bound to the pseudo-variables that label the arcs of the path. In our example, these expressions are less-than inequalities. If the value of the program variable on the left of the inequality grows to reach the expressed bound after a finite number of repetitions of the path then the corresponding pseudo-variable will eventually become 0. Sufficient conditions for this to occur are that the program variable is (1) monotonic with respect to executions of the transition function corresponding to the path and (2) increasing almost everywhere in any sequence of repetitions of the path.

Consider the path from the first occurrence of a node labeled W to its repeated occurrence, controlled by pseudo-variable x_2 . The program variable that appears in the corresponding inequality is *state.time*. The path controlled by x_2 represents only the transitions S_{warn} and S_{tick} , which are enabled by




(a)

(b)



(c)



(d)

Figure 3

external events *Ready* and *Tick*, respectively. Inspection of the coalgebra specification *Reflex* reveals that a *\$warn* transition does not change *state.time*, while a *\$tick* transition increments its value. Thus the path satisfies condition (1), monotonicity of the program variable. Furthermore, since the *Tick* event occurs almost everywhere in an external event stream, the *\$tick* transition will occur almost everywhere in a sequence of transitions from state *W* to itself, thus the path also satisfies condition (2). This argument proves the temporal assertion $\diamond(x_2 = 0)$.

This allows the BDD of Figure 3(a) to be reduced by removing the repeated occurrence of node *W* and the arc leading to it. Similar reasoning justifies removal of the repeated occurrence of nodes *R*, *E*, and the first repeated occurrence of node *S* leaving the BDD depicted in Figure 3(b). The paths not controlled by a pseudo-variable in this BDD are not of interest and can be reduced to single arcs.

In Figure 3(c), the arcs have been labelled with sets of the pairs of event sequences and corresponding transition actions represented by an interpretation of the arc in the *Reflex* coalgebra. The set of transition sequences from the root node to its repeated occurrence is gotten by taking the cartesian product of the sequence of set-valued labels. As the labels on individual arcs are all singletons, so is the composite label, which is $\{([Ready, Stop], [Sready, $go, $react, $tock])\}$.

We examine the expressions x_1 , x_3 and x_5 in the context of the path transition sequence. The program variable *state.time* is not monotonic with respect to the transition sequence, hence we cannot conclude that repeated extensions of the path would cause the variables x_1 or x_3 to assume zero values. However, the program variable *state.trial.number* is monotonically increasing over this transition sequence, thus the arc labeled by x_5 and the repeated node *S* can be eliminated. The resulting BDD, depicted in Figure 3(d), is reducible to the singleton node **1** as its canonical form. This constitutes a proof of the conjectured liveness property.

4 CONCLUSIONS

We have introduced a functional programming notation that does not depend upon explicit recursion in definitions but uses instead the structure of signature coalgebras. The important contributions of this notation and the mathematical structures that underlie it are:

- It provides a framework in which control and data transformation are separately specified. A program is specified by the transitions of a finite-state system.
- All familiar iteration schemes can be modeled by varieties of coalgebras.
- Each variety of coalgebra has associated with it proof rules that virtually dictate the form of proofs of safety properties of algorithms that conform to its structure.

- Liveness properties can be verified through a hybrid deduction scheme in which temporal logical inference is used in conjunction with symbolic model checking.

REFERENCES

- [1] Samson Abramsky. A domain equation for bisimulation. *Information and Computation*, 92:161–218, 1992.
- [2] Edward A. Ashcroft and William W. Wadge. Lucid: A non-procedural language with iteration. *Communications of the ACM*, 20(7):519–526, 1977.
- [3] G. Berry and G. Gonthier. The Esterel synchronous programming language: Design, semantics, implementation. *Science Of Computer Programming*, 19(2):87–152, 1992.
- [4] Richard S. Bird. An introduction to the theory of lists. In M. Broy, editor, *Logic of Programming and Calculi of Discrete Design*, volume 36 of *NATO Series F*. Springer-Verlag, 1986.
- [5] G. Boudol, V. Roy, R. de Simone, and D. Vergamini. Process calculi, from theory to practice: Verification tools. In *Automatic Verification Methods for Finite State Systems, LNCS 407*, pages 1–10. Springer-Verlag, 1990.
- [6] R. E. Bryant. Symbolic boolean manipulation with ordered binary decision diagrams. *ACM Computing Surveys*, 24(3):293–318, September 1992.
- [7] R. M. Burstall and P. J. Landin. Programs and their proofs: an algebraic approach. In *Machine Intelligence 4*, pages 17–44. Edinburgh University Press, 1969.
- [8] Rod Burstall. Inductively defined functions in functional programming languages. *Journal of Computer and System Sciences*, 34(2/3):409–421, 1987.
- [9] P. Caspi and M. Pouzet. Synchronous Kahn networks. In *Proceedings of ACM SIGPLAN Internat. Conf. on Functional Programming*, pages 226–238. ACM Press, May 1996.
- [10] E. M. Clarke, O. Grumberg, and D. Long. Verification tools for finite-state concurrent systems. In *A Decade of Concurrency: Reflections and Perspectives*, volume 803 of *Lecture Notes in Computer Science*, pages 124–175. Springer Verlag, 1992.
- [11] E. A. Emerson and E. M. Clarke. Using branching time temporal logic to synthesize synchronizations skeletons. *Science of Computer Programming*, 2:241–266, 1982.
- [12] Marcelo P. Fiore. A coinduction principle for recursive data types based on bisimulation. *Information and Computation*, 127(2):186–198, 1996.
- [13] Peter Freyd. Recursive types reduced to inductive types. In *Fifth IEEE Symposium on Logic in Computer Science*, pages 498–507. IEEE Computer Society Press, 1990.
- [14] Andrew Gordon. Bilimilarity as a theory of functional programming. In *Proceedings of 11th Conference on Mathematical Foundations of Programming Semantics*, volume 1 of *Electronic Notes in Theoretical Computer Science*. Elsevier, URL: www.pigeon.elsevier.nl/mcs/tcs/pc/volume1.htm, 1995.
- [15] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud. The synchronous dataflow programming language Lustre. *Proceedings of the IEEE*, 79(9):1305–1320,

- 1991.
- [16] N. Halbwachs, J.-C. Fernandez, and A. Bouajjanni. An executable temporal logic to express safety properties and its connection with the language Lustre. In *Sixth International Symp. on Lucid and Intensional Programming, ISLIP'93*, Quebec City, Canada, April 1993. Université Laval.
 - [17] N. Halbwachs, F. Lagnier, and C. Ratel. Programming and verifying real-time systems by means of the synchronous data-flow programming language lustre. *IEEE Transactions on Software Engineering*, 18(9), 1992.
 - [18] Bart Jacobs. Mongruences and cofree coalgebras. In *Algebraic Methodology and Software Technology – AMAST'95*, volume 936 of *Lecture Notes in Computer Science*, pages 245–260. Springer Verlag, 1995.
 - [19] Bart Jacobs and Jan Rutten. A tutorial on (co)algebras and (co)induction. *EACTS Bulletin*, 62:222–259, 1997.
 - [20] Gilles Kahn. The semantics of a simple language for parallel programming. In *IFIP'74 Congress*. North Holland, 1974.
 - [21] Richard B. Kieburtz and Jeffrey Lewis. Algebraic Design Language—Preliminary definition. Technical report, Pacific Software Research Center, Oregon Graduate Institute of Science & Technology, January 1994.
 - [22] Richard B. Kieburtz and Jeffrey Lewis. Programming with algebras. In *Advanced Functional Programming*, volume 925 of *Lecture Notes in Computer Science*, pages 267–307. Springer Verlag, 1995.
 - [23] Erik Meijer, Maarten Fokkinga, and Ross Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In *Proc. of 5th ACM Conf. on Functional Programming Languages and Computer Architecture*, volume 523 of *Lecture Notes in Computer Science*, pages 124–144. Springer-Verlag, August 1991.
 - [24] Robin Milner and Mads Tofte. Co-induction in relational semantics. *Theoretical Computer Science*, 87:209–220, 1992.
 - [25] Lawrence C. Paulson. Mechanizing coinduction and corecursion in higher-order logic. *Journal of Logic and Computation*, 7(2):175–204, 1997.
 - [26] Andrew Pitts. A coinduction principle for recursively defined domains. *Theoretical Computer Science*, 124(2):195–219, 1994.

About the Author

Richard B. Kieburtz is Professor of Computer Science and Engineering at the Oregon Graduate Institute. His research interests are in the areas of functional programming; formal methods for specification and construction of software; verification and its application to hardware and software systems.