Syntax and Semantics of the SpecC Language

Jianwen Zhu, Rainer Dömer, Daniel D. Gajski
Department of Information and Computer Science
University of California, Irvine, CA 92697-3425, USA
{jzhu, doemer, gajski}@ics.uci.edu

In this paper, we describe an executable modeling language in the context of a homogeneous codesign methodology featuring synthesis, reuse and validation. A C based language called SpecC is proposed as an attempt to achieve these goals. The syntax and semantics of the language is presented and compared with existing HDLs. In conclusion, SpecC is conceptually more abstract, syntactically simpler, and semantically richer.

I. Introduction

To tackle the complexities involved in systems-on-a-chip design, designers often follow two strategies. In favor of synthesis, the top-down strategy starts from an abstract specification of the systems functionality, and then performs stepwise refinement into implementations with the help of various synthesis tools. In favor of reuse, the bottom-up approach implements the systems functionality by integrating existing components. In reality, these two approaches have to be combined, because for the top-down approach, it is often unaffordable for new designs to start from scratch, whereas for the bottom-up approach, it is rarely possible to build systems completely from existing components without adding new ones.

New codesign methodologies have to be developed to support the seamless mixture of both strategies. One such methodology is illustrated in Fig.1(a), where boxes represent **design tasks**, and ellipses represent design descriptions, or their abstract form, the **design models**, which are shown in more detail in Fig.1(b).

For the synthesis flow, designers start from the specification stage, where the intended functionality of the system is specified in terms of a set of behaviors communicating via a set of channels. Note that at this stage, the designers may already decide to use some components from the reuse library, so these components should also be able to be instantiated in the specification.

At the second stage, designers partition the behaviors onto ASICs and programmable processors. The design model generated at this stage should carry these design decisions.

At the communication and synchronization refinement stage, the abstract channels among behaviors have to be resolved into implementations, such as behaviors transferring data over a bus according to a particular protocol.

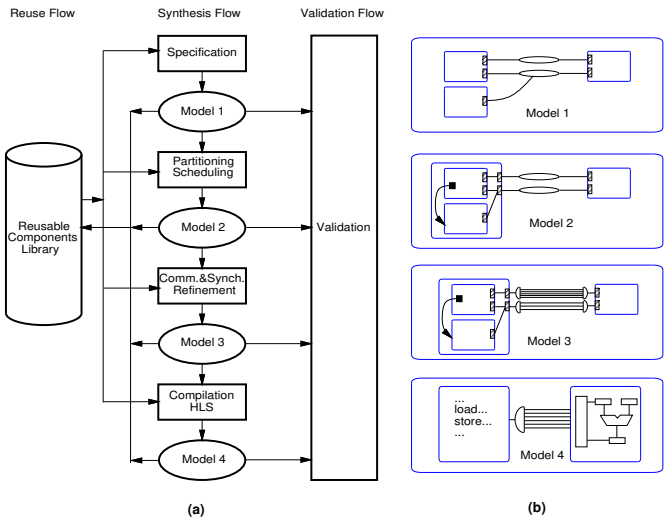

Fig. 1. A homogeneous codesign methodology

The design model generated at this stage simply replaces the abstract channels by detailed ones.

The final stage might be the compilation of software into binary code and high level synthesis of ASIC behavior into register transfer level designs. The model generated at this stage might be a software behavior (in the form of an instruction stream) communicating with a hardware behavior (in the form of a RTL netlist) via a bus functional model, which translates each instruction into a series of activities on the bus wires.

For the validation flow, the design model generated at each stage should be simulated before proceeding to the next synthesis step. For the reuse flow, there exists a library of reusable components, which as a repository of design models, can be "downloaded" or "uploaded" by synthesis tools and designers.

For this methodology, it is desirable that we

Objective 1: capture the design model at each design stage in a **single language**;

Objective 2: make the language **executable**.

Objective 3: encourage **reuse**;

Objective 4: facilitate **synthesis**.

Meeting of Objective 1 will greatly simplify the code-

sign methodology since synthesis tasks are now merely transformations from one program into another using the same language, and validation tasks are just the program executions. This **homogeneous** approach is in contrast with traditional approaches where the systems functionality is specified using one specification language, or a mixture of several, and is then opaquely synthesized into design descriptions represented by different languages, for example, C for software, VHDL for hardware.

The requirement of executability (Objective 2) of the language is of crucial importance for validation. For example, the specification needs to be validated to make sure that the functionality captured is what the designer intended. Validation is even necessary for the intermediate design models, for example those after the communication refinement stage, where the interaction between hardware and software becomes very tricky and error-prone.

For reuse (Objective 3), it is important that the design models stored in the library are captured in the same language, so that they can be checked in and out easily. Furthermore, the language should be designed in such a way that a design model of one component is decoupled from those of other components.

In order to facilitate synthesis (Objective 4) incremental refinement is used. Since we use the same language to model the design at each stage, a synthesis task is a refinement of one program into a more detailed program. The change made by the transformation should be local, meaning that it will only change a local part of the program without affecting other parts. For example, the communication refinement task should only replace abstract channels into more detailed ones without changing the model of the behaviors which use these channels. The locality of changes makes the synthesizer simpler and the generated results more understandable.

In this paper, we attempt to achieve the objectives above by proposing a new C based language called SpecC. SpecC supports all concepts commonly found in embedded systems, such as concurrency, state transitions, structural and behavioral hierarchy, exception handling, timing, communication and synchronization, as discussed in [6]. These concepts have to be organized in an orthogonal way so that the language can be minimal. Section II. presents the set of constructs in SpecC that support these concepts. While Section II. shows that SpecC is conceptually more abstract and syntactically simpler than VHDL, Section III.describes why it is semantically richer. Section IV. shows how the objectives defined above are achieved by presenting an example system. Finally Section V. concludes this paper with comparisons to traditional HDLs.

II. Syntax

In this section, we give an introduction of the constructs provided by SpecC.

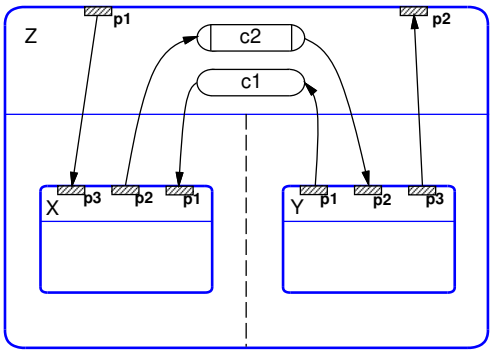

(a)

```
interface ILeft {
      void   write( int val );
      };
interface IRight {
      int    read( void );
      };

channel CSharedI( void )
        implements ILeft, IRight {
      int    storage;
      bool   valid;

      void   write( int val ) {
             storage = val;
             valid = true;
             }
      int    read( void ) {
             while( !valid );
             return storage;
             }
      };

behavior  Y( out int p1, IRight p2, out int p3 ) {
      int    local;

      void   main( void ) {
             ...
             local = p2.read();
             ...
             }
      };
```

```
behavior  X( in int p1, ILeft p2, in int p3 ) {
      int    local;

      void   main( void ) {
             ....
             p2.write( local );
             ...
             }
      };
behavior  Z( in int p1, out int p2 ) {
      int          c1;
      CShared   c2;

      X      x( c1, c2, p1 );
      Y      y( c1, c2, p2 );

      void   main( void ) {
             par { x.main(); y.main(); }
             }
      };
```

(b)

Fig. 2. Basic structure of a SpecC program

A. Basic structure

The SpecC view of the world is a hierarchical network of **behaviors** interconnected by **channels**.

Each behavior possesses

- a set of ports, such as *p1* and *p2* of behavior *Z* in Fig.2;
- a set of state variables;
- a set of channels, such as *c1* and *c2* in behavior *Z* in Fig.2;
- a functionality, which contains computations updating the state variables and communicating via the ports.

The behavior may be a composite behavior which contains a set of child behaviors, in which case the functionality of the composite behavior is specified by composing its child behaviors. For example, in Fig.2, the functionality of behavior *Z* is composed of the parallel execution of behaviors *X* and *Y*.

SpecC provides the *behavior* construct to capture all the information for a behavior. This *behavior* construct looks like a C++ class which exports a *main* method. Ports are

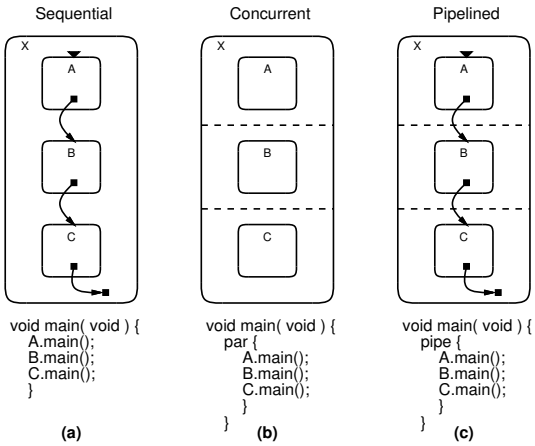

Fig. 3. Behavioral hierarchy

declared in the parameter list. State variable, channels and child behavior instances are declared as typed variables, and the functionality is specified by the methods, or functions starting from *main*. The *behavior* construct can be used as a type to instantiate behavior instances.

A channel is an entity responsible for communication. It can be a primitive channel in the form of built-in data types such as *int, char* and *float*, or it can be a complex channel whose specification is split into two constructs: the *interface* construct declares what kind of communications a channel can perform; the *channel* construct provides the implementation of how the communication is performed. The *interface* is usually used as a type to declare ports, whereas the *channel* is used as a type to instantiate channels. Section C. explains in detail how complex channels are specified and why they are useful.

In summary, the SpecC program consists of a list of specifications of behaviors and channels.

B. Hierarchy

The *behavior* construct can orthogonally capture both **structural** hierarchy and **behavioral** hierarchy.

Structural hierarchy means that composite behaviors can be decomposed into child behaviors interconnected by channels. Channels define the set of paths through which the child behaviors communicate, and how the communication is performed. When the communication is performed is determined by the functionality of the child behaviors. Structural information is captured by behavior instantiation where its ports are mapped to channels or the ports of the parent behavior.

Behavioral hierarchy means that the functionality of a composite behavior is composed of the functionality of the child behaviors. There are three constructs to specify how child behaviors are composed in time into a composite behavior. For example, we can specify a behavior being the sequential composition of the child behaviors using sequential statements, as shown in Fig.3(a), where X finishes when the last behavior C finishes. Second, we

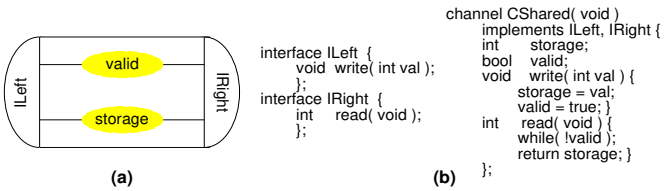

Fig. 4. Shared memory channel

can use the parallel composition using the *par* construct, as shown in Fig.3(b), where X finishes when all its child behaviors A, B and C are finished. Also, pipelined composition is supported using the *pipe* construct, as shown in Fig.3(c), where X starts again when the slowest behavior finishes.

In summary, structural hierarchy is captured by the tree of behavior instantiations, whereas behavioral hierarchy is captured by the tree of function calls to the behavior *main* methods.

C. Communication

We have mentioned in Section A. that the channel concept is used to model communication and the specification of a complex channel is split into the *interface* and *channel* constructs. The *interface* construct encapsulates a set of method prototype declarations, which specify what kind of communications a channel can perform. The channel encapsulates a set of media in the form of variables, for example a set of storages or a set of wires, and the set of method implementations which specify how the communication is performed.

Fig.4 shows a shared variable channel which can be accessed by concurrent behaviors. The communications that can be performed are *read* and *write*, as declared in the *ILeft* interface and *IRight* interface respectively.

The channel *CShared* encapsulates the variable *storage* being shared, a *valid* bit, which the *write* operation has to set and the *read* operation has to spinwait for, as well as definitions of the *read* and *write* methods.

The channel is related to the interfaces by the *implements* keyword followed by the list of interfaces, which implies that it is mandatory for the channel to implement the methods declared in the interfaces.

The difference between methods in a channel and methods in a behavior has to be emphasized. While methods of a behavior specify the functionality of itself, methods of a channel specify functionality of the caller, in other words, they will get **inlined** into the connected behaviors when the system is finally implemented. When a channel is inlined, the encapsulated media are exposed, the methods are moved to the caller and the interface port is flattened into ports connected to the exposed media.

The fact that a port of an interface type can be resolved to a real channel at the time of its instantiation is called **late binding**. Such a late binding mechanism makes it possible for a behavior to perform function calls via the

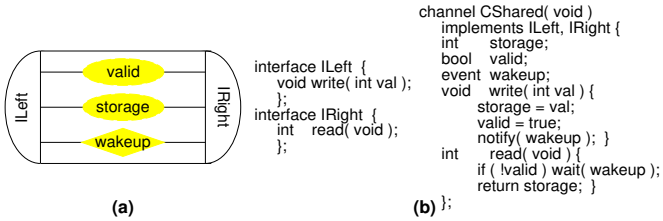

Fig. 5. Event and shared memory channel

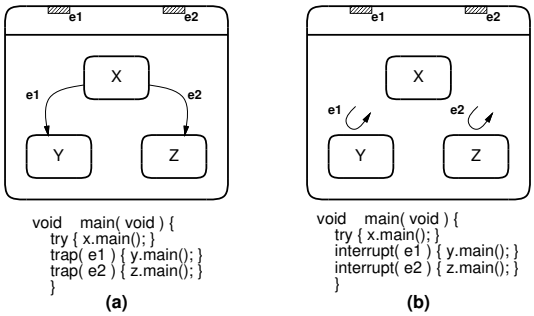

Fig. 6. Exception handling: (a) abortion, (b) interrupt.

interface port without knowing what kind of channel it will be eventually mapped to. In this way, any channel can be plugged in as long as it conforms to this interface.

Such a "plug-and-play" feature is essential for both reuse and incremental refinement. For example, it is possible to replace the channel *CShared* in Fig.2(b) with another, more detailed channel, or one that uses a different protocol, without affecting the description of the behaviors *X, Y* and *Z*, as long as that channel implements both interfaces *ILeft* and *IRight*.

D. Synchronization

Concurrent behaviors often need to be synchronized in order to be cooperative. In SpecC a built-in type *event* serves as the basic unit of synchronization. An event can be considered as a channel with a set of methods including *wait* and *notify*. A *wait* call on an event will suspend the caller. A *notify* call on an event will resume all behaviors that are suspended due to a wait on this event.

The shared memory example introduced in Fig.4 can be rewritten using the event mechanism as shown in Fig.5. In Fig.4(b), the *read* method has to poll the valid bit constantly, which is unefficient. A better way, as shown in Fig.5(b), is to use an event called *wakeup*, so that a call to the *read* method can suspend itself when *valid* is false, and will be resumed later when a write operation notifies the *wakeup* event.

E. Exceptions

In order to model exceptions in an embedded system SpecC supports two concepts, namely **abortion** and **interrupt**, as shown in Fig.6.

The *try-trap* construct shown in Fig.6(a) aborts behavior x immediately when one of the events $e1$, $e2$ occurs. The execution of behavior x (and all its child behaviors) is terminated without completing its computation and control is transferred to behavior y in case of $e1$, to behavior z in case of $e2$. This type of exception usually is used to model the reset of a system.

On the other hand the *try-interrupt* construct, as shown in Fig.6(b), can be used to model interrupts. Here again execution of behavior x is stopped immediately for events $e1$ and $e2$, and behavior y or z, respectively, is started to service the interrupt. After completion of interrupt handlers y and z control is transferred back to behavior x and execution is resumed right at the point where it was stopped.

For both types of exceptions, in case two or more events happen at the same time, priority is given to the first listed event.

F. Timing

In the design of embedded systems the notion of real time is an important issue. However, in traditional imperative languages such as C, only the ordering among statements is specified, the exact information on when these statements are executed, is irrelevant. While these languages are suitable for specifying functionality, they are insufficient in modeling embedded systems because of the lack of timing information. Hardware description languages such as VHDL overcome this problem by introducing the notion of time: statements are executed at discrete points in time and their execution delay is zero. While VHDL gives an exact definition of timing for each statement, such a treatment often leads to **over-specification**. One obvious over-specification is the case when VHDL is used to specify functional behavior. The timing of functional behaviors is unknown until they are synthesized. The assumption of zero execution time is too optimistic and there are chances to miss design errors during specification validation.

SpecC overcomes this problem by differentiating between **timed behavior**, which executes in zero time, and **untimed behavior**, on which no assumption of timing can be made. The execution semantics of SpecC as described in Section III. will allow the synchronized execution of timed and untimed behavior.

Other cases of over-specification are timing constraints and timing delays, where events have to happen, or, are guaranteed to happen in a **time range**, instead of at a fixed point in time, as restricted by VHDL.

SpecC supports the specification of timing explicitly and distinguishes two types of timing specifications, namely **constraints** and **delays**. At the specification level timing constraints are used to specify time limits that have to be satisfied. At the implementation level computational delays have to be noted.

Consider, for example, the timing diagram of the read protocol for a static RAM, as shown in Fig.7(a). In order to read a word from the SRAM, the address of the data

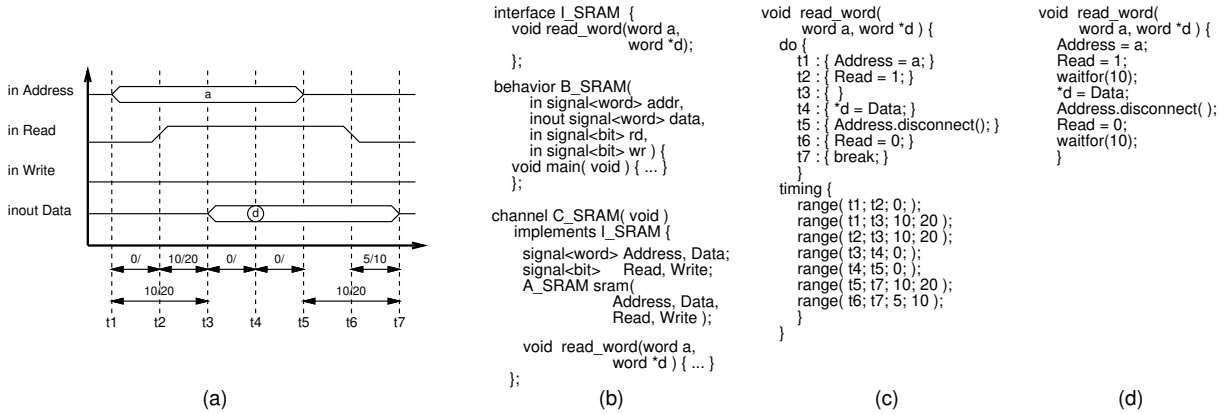

Fig. 7. SRAM read protocol: (a) timing diagram, (b) SRAM channel, (c) specification level timing, (d) implementation level timing.

is supplied at the *address* port and the read operation is selected by assigning 1 to the *read* and 0 to the *write* port. The selected word then can be accessed at the *data* port. The diagram in Fig.7(a) explicitly specifies all timing constraints that have to be satisfied during this read access. These constraints are specified as arcs between pairs of events annotated with x/y, where x specifies the minimum and y the maximum time between the value changes of the signals. The times are measured in real time units such as nanoseconds.

Fig.7(b) shows the SpecC source code of a SRAM channel *C_SRAM*, which instantiates the behavior *B_SRAM*, and the signals, which are mapped to the ports of the SRAM. Access to the memory is provided by the *read_word* method, which encapsulates the read protocol explained above (due to space constraints write access is ignored).

Fig.7(c) shows the source code of the *read_word* method at the specification level. The *do-timing* construct used here effectively describes all information contained in the timing diagram. The first part of the construct lists all the events of the diagram, which are specified as a label and its associated piece of code, which describes the changes of signal values. The second part is a list of *range* statements, which specify the timing constraints using 4-tuples $T = (e1, e2, min, max)$, where $e1$ and $e2$ are event labels and min and max specify the minimum and maximum time, respectively, between these events.

This style of timing description is used at the specification level. In order to get an executable model of the protocol **scheduling** has to be performed for each *do-timing* statement. Fig.7(d) shows the implementation of the *read_word* method after an ASAP scheduling is performed. All timing constraints are replaced by delays, which are specified using the *waitfor* construct.

III. Execution semantics

In this section, we give a formal definition of the execution semantics of a SpecC program. We use a graph-based notation similar to [4].

A. Graph-based representation

In its most primitive form with all the high level constructs flattened and syntactical sugar stripped, the SpecC program can be represented by a directed graph. Fig.8 shows an example of such a graph.

The flattened SpecC program can be broken into blocks of statement sequences, which become vertices of the graph. Four types of constructs, which become edges of the graph, break the program into blocks:

Type 1: branches, including *if-else*, *while*, *for* loops etc,

Type 2: delay (*waitfor*) statements,

Type 3: *wait* statements,

Type 4: *notify* statements.

Note that while Type 1 constructs break the program into basic blocks as in traditional control flow analysis, Type 2-4 constructs break the basic blocks further.

A vertex is called a **timed** vertex, if the execution of the vertex takes zero time, as determined by the syntax of the language. It is referred to as an **untimed** vertex otherwise. Another way of classifying the vertices is to call a vertex a **reactive** vertex if it is preceded by a *wait* statement. Otherwise, it is called a **non-reactive** vertex.

Furthermore, there are special vertices used to represent high level language constructs such as *par, cpar* and *pipe*, namely, *fork* and *join* vertices. For the simplicity of the presentation, we ignore the constructs *try-trap* and *try-interrupt* in this discussion.

There are two types of edges. The **sensitizing** edges, represented by dotted arrows in Fig.8(b), are derived from the sequencing information of the program where the control flows to reactive vertices. The execution of the source vertex of a sensitizing edge will make the sink vertex sensitive to the events that it waits for. Each sensitizing edge is associated with a *condition*.

The **triggering** edges are represented by solid arrows in Fig.8(b) and are annotated by *condition/delay* pairs.

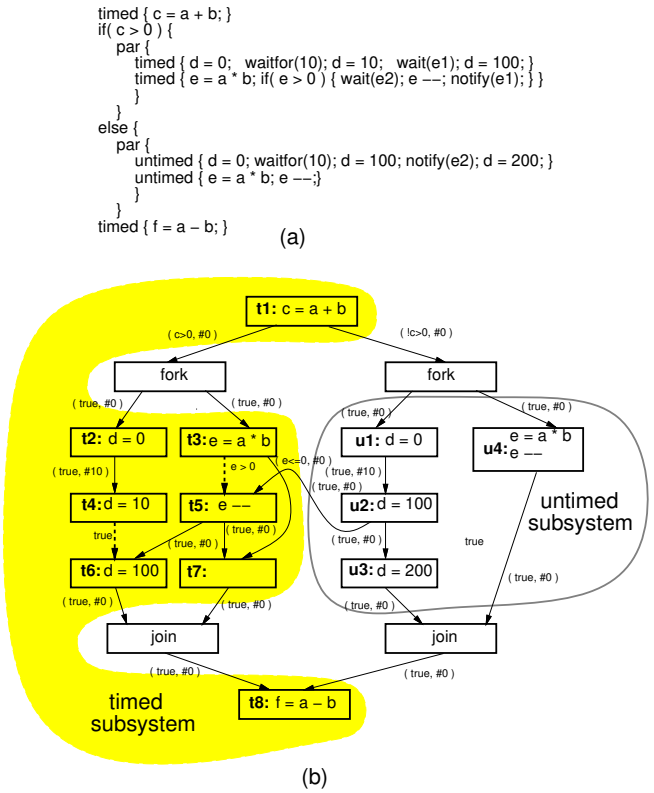

Fig. 8. Execution semantics: (a) pseudo SpecC code, (b) graph.

The execution of the source vertex of a triggering edge will cause the sink vertex to execute at *delay* time steps later if the *condition* is true. If the sink vertex is a reactive vertex, it also has to be sensitive in order to be executed. There are three sources of program information for the triggering edges: (1) sequencing information where control flows to non-reactive vertices; (2) delay information associated with *waitfor* statements; (3) synchronization information associated with *wait-notify* pairs.

The subgraph spanned by the set of timed vertices is called the **timed subsystem**. The subgraph spanned by the untimed vertex set is called the **untimed subsystem**. The timed subsystem is intended to model the synthesized hardware using discrete event semantics equivalent to VHDL. The untimed subsystem is intended to model unsynthesized behavior or software, whose timing is unknown.

More formally, a SpecC program can be represented as a graph $G = \langle V, v_0, E_s, E_t \rangle$, where

- $V = T \cup U \cup Fork \cup Join$, where T is the set of timed vertices, U is the set of untimed vertices, $Fork$ is the set of *fork* vertices, $Join$ is the set of *join* vertices;
- $v_0 \in V$ is the start vertex;
- $E_s \subseteq V \times V \times B$ represents the conditional sensitizing edges, where B is the set of boolean expressions;
- $E_t \subseteq V \times V \times B \times Z^+$ represents the conditional triggering edges with delays, where Z^+ is the set of positive integers.

For convenience of notation, we define the set of reactive vertices to be $W = \{v | \exists v' \in V, b \in B, \langle v', v, b \rangle \in Es\}$.

B. Semantics

The state of the computation can be represented by $s = \langle N, R, D, M, C \rangle$, where

- $N \subseteq V$ is the set of sensitive vertices,
- $R \subseteq V$ is the set of vertices ready to be executed,
- $D \subseteq V \times Z^+$ is the set of all pairs $\langle v, d \rangle$, such that vertex v has d time steps left before it can execute,
- M is the memory store that maps each variable to its current value,
- C is the memory store which holds the execution context of all the vertices.

We assume the existence of the following operations:

- $v = SelectOneOf(R)$ deterministically chooses one vertex v from the set of ready vertices R .
- $M' = ExecTimed(v, M)$ applies the code associated with timed vertex v to the memory store and returns a new memory store M' .
- $\langle t, M', C' \rangle = ExecUntimed(v, M, C)$ applies the code associated with untimed vertex v under a context C to the memory store for some host time. It will return a triple $\langle t, M', C' \rangle$, where $t \in Z^+$ represents the time steps elapsed during the execution, M' represents the updated memory store and C' represents the new context. Note that the granularity for each execution of the untimed vertex, that is, how much code it executes and how many time steps it will take, can be specified by the user.
- $Finished(v, C)$ returns whether vertex v has finished its execution, where C is the context.
- $Eval(b, M)$ evaluates the boolean expression b using the memory store M and returns TRUE or FALSE.

We define $s' = Next(\langle N, R, D, M, C \rangle)$ as $\langle N', R', D', M', C' \rangle$, where

- **Type I Transition** represents the **execution** of a vertex at the current time step:
  if $R \neq \emptyset$, then Let $v = SelectOneOf(R)$
  if $v \in T$, then

$$\begin{aligned} M' &= ExecTimed(v, M) \\ C' &= C \\ N' &= N - \{v\} \cup \{v' | \langle v, v', b \rangle \in Es \\ &\quad \wedge Eval(b, M')\} \\ R' &= R - \{v\} \cup \{v' | \langle v, v', b, d \rangle \in Et \\ &\quad \wedge Eval(b, M') \wedge d = 0 \\ &\quad \wedge (v \in W \wedge v \in N' \vee v \notin W)\} \\ D' &= D \cup \{\langle v', d \rangle | \langle v, v', b, d \rangle \in Et \\ &\quad \wedge Eval(b, M') \wedge d > 0\} \end{aligned}$$

  else if $v \in Fork$, or if $v \in Join \wedge$

$$\forall v' \in \{v' | \langle v', v, b \rangle \in Et\},\ Finished(v', C)$$

then

$$\begin{aligned} M' &= M \\ C' &= C \\ N' &= N - \{v\} \cup \{v' | \langle v, v', b\rangle \in Es \\ &\quad \wedge\ Eval(b, M')\} \\ R' &= R - \{v\} \cup \{v' | \langle v, v', b, d\rangle \in Et \\ &\quad \wedge\ Eval(b, M')\ \wedge\ d = 0 \\ &\quad \wedge\ (v \in W\ \wedge\ v \in N' \vee v \notin W)\} \\ D' &= D \cup \{\langle v', d\rangle | \langle v, v', b, d\rangle \in Et \\ &\quad \wedge\ Eval(b, M')\ \wedge\ d > 0\} \end{aligned}$$

else if $v \in Join \wedge$

$$\exists v' \in \{v' | \langle v', v, b\rangle \in Et\},\ \neg Finished(v', C)$$

then

$$\begin{aligned} M' &= M, \\ C' &= C \\ N' &= N \\ R' &= R - \{v\} \\ D' &= D \end{aligned}$$

else if $v \in U$, then

$$\langle t, M', C'\rangle = ExecUntimed(v, M, C)$$

if $Finished(v, C')$ then

$$\begin{aligned} N' &= N - \{v\} \cup \{v' | \langle v, v', b\rangle \in Es \\ &\quad \wedge\ Eval(b, M')\} \\ R' &= R - \{v\} \cup \{v' | \langle v, v', b, d\rangle \in Et \\ &\quad \wedge\ Eval(b, M')\ \wedge\ d = 0 \\ &\quad \wedge\ (v \in W\ \wedge\ v \in N' \vee v \notin W)\} \\ D' &= D \cup \{\langle v', d + t\rangle | \langle v, v', b, d\rangle \in Et \\ &\quad \wedge\ Eval(b, M')\ \wedge\ d > 0\} \end{aligned}$$

else

$$\begin{aligned} N' &= N \\ R' &= R - \{v\} \\ D' &= D \cup \{\langle v, t\rangle\} \end{aligned}$$

- **Type II Transition** represents the **advance of time** to the next step when a vertex can execute.
  if $R = \emptyset$ and $D \neq \emptyset$, then let

$$\begin{aligned} d_0 &= min_{\langle v,d\rangle \in D} d \\ A &= \{\langle v, d\rangle | \langle v, d\rangle \in D\ \wedge\ d = d_0\} \end{aligned}$$

  then

$$\begin{aligned} M' &= M \\ N' &= N \\ R' &= \{v | \langle v, d\rangle \in A \\ &\quad \wedge\ (v \in W\ \wedge\ v \in N' \vee v \notin W)\} \\ D' &= \{\langle v, d\rangle | \langle v, d + d_0\rangle \in D - A\} \\ C' &= C \end{aligned}$$

- **Type III Transition** represents the **end of execution**: if $R = \emptyset$ and $D = \emptyset$, then $s = \perp$, denoting no next state.

The execution of graph G is a sequence of states $[s_0, s_1, \ldots, \perp]$, where

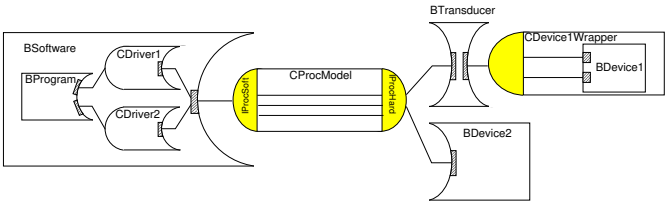

Fig. 9. Model of a microprocessor system with IO devices.

- $s_0 = \langle \emptyset, \{v_0\}, \emptyset, M_0, C_0\rangle$, where M_0 represents the initial memory store that maps every variable to its initial value, and C_0 represents the initial context which sets all the untimed vertices to its first statement.
- $s_{i+1} = Next(s_i, G)$

In summary, we present the semantics of the SpecC language in a formal graph-based notation. The semantics is rich in the sense that it covers the semantics of many other languages. For example, if we constrain all the triggering edges to a delay of zero (disallow *waitfor* statements), then the triggering edges induce a partial ordering on the vertices, and there is no notion of the passage of time. Such a system is equivalent to the systems captured by concurrent languages. If we further constrain that no *fork* and *join* vertices are allowed, such a system is equivalent to those captured by sequential imperative languages such as C. On the other hand, if we exclude the untimed subsystem, the language is semantically equivalent to VHDL.

IV. Meeting the objectives

In this section, we briefly review how the SpecC language addresses the four design objectives discussed in Section I. and summarize SpecC features.

Objective 1 requires SpecC to be capable of modeling designs at different abstraction levels, or mixed levels of abstractions. In the codesign domain, a **computation** at a high abstraction level may be a behavior with only partial ordering of operations specified, but exact timing missing. On the other hand, a computation with lower abstraction level may be a behavior with the exact information when each operation is performed. Similarly, a **communication** at a high abstraction level may be a shared variable accessable by concurrent processes. At a lower abstraction level, these shared variables may be distributed over different processing elements, while accesses to them may involve consistency protocols and complex transactions over system buses.

For computation, SpecC allows the simultaneous specification of an untimed system, which is primarily used to model unsynthesized functionality whose timing is not resolved yet, and a timed system, which is primarily used to model synthesized functionality with known timing. This flexibility makes it possible to describe a variety of system modeling configurations. For example, Fig.9 shows a typical microprocessor system with two IO devices. Software

behavior, which is not yet compiled and bound to any processor, can be modeled as untimed behavior, such as behavior *BProgram* in Fig.9. So does a hardware behavior that is to be synthesized, such as behavior *BDevice2*. Compiled software behavior can be modeled as timed behavior, such as channels *CDriver1* and *CDriver2* representing the device drivers for behaviors *BDevice1* and *BDevice2*. A processor model *CProcModel* is modeled as a channel which exports its instruction set in the form of methods. Each instruction is modeled as timed behavior, which operates on the processor bus signals. Hardware components from the library, that is, synthesized hardware behaviors, can be modeled as timed behavior, for example, behavior *BDevice2* in Fig.9.

For communication, SpecC allows the simultaneous specification of communication using primitive channels, which is used to model abstract communication via implicit read and write operations over typed variables, and complex channels, which is primarily used to model communication at the implementation level, for example, the transfer of a block of data over a standard bus.

The execution semantics of SpecC laid the basis to realize Objective 2. The benefit of being able to execute models at mixed abstraction levels is two fold: first, intermediate design models can be validated before the next synthesis step; second, designs can be validated with appropriate speed-accuracy trade-off. For example, in Fig.9, behavior *BProgram* is modeled as untimed behavior, while channels *CDriver1* and *CDriver2* are modeled as timed behaviors. The rationale behind this configuration is that the code in *CDriver1* and *CDriver2* contains IO instructions which interact intensively with the hardware such as *BDevice2* and *BTransducer*, and it is this type of code that is error-prone and should receive our attention. On the other hand, the code contained in behavior *BProgram* contains just normal operations of the processor and does not need to be verified at such a detailed level.

The ability of SpecC to model designs at different abstraction levels makes a large category of design artifacts eligible to be modeled and entered in a reuse library (Objective 3). For example, a hardware component can be modeled as a behavior, such as *BDevice1* in Fig.9, and can be stored in the library. The communication protocol for that hardware component, which before is documented by a data sheet, can now be modeled as a wrapper, such as *CDevice1Wrapper* in Fig.9, and also is stored in the library. The bus protocols, including standard system buses such as *PCI bus*, or *VMEBus*, and processor bus models, such as *Pentium* bus and *PowerPC* bus, can also be stored in the library in the form of channels.

The abstraction of communication into a set of functions in the *channel* construct and the abstraction of the channel implementation into a set of function prototypes in the *interface* construct makes it possible to decouple the computational aspect of a behavior from the communication. This feature is helpful for realizing both Objective 3 and Objective 4. For reuse, the behaviors described in this way can be used without modification in different situations. For synthesis and refinement, abstract channels can be replaced by detailed channels without affecting the connected behaviors.

V. Conclusion

In conclusion, we proposed SpecC as a modeling language for codesign, which supports a homogeneous codesign methodology. SpecC can be considered as an improvement over traditional HDLs such as VHDL.

Semantically, SpecC allows the specification of behavior with exact timing as well as unknown timing, whereas VHDL only allows specification of behavior with exact timing, which often leads to overspecification.

Conceptually, SpecC raises the abstraction level, while reorganizing important concepts. For example, concurrency is decoupled from structure, synchronization and timing are decoupled from interconnections (for example, in VHDL signals are used both for synchronization, interconnection, and even timing), function interfaces are decoupled from function implementations, ports are generalized into interfaces, wires are generalized into channels.

Syntactically, SpecC is based on C, which allows the inheritance of a large archive of existing code, and makes it easy for an implementation of the language to leverage traditional C compilers.

Philosophically, SpecC is intended to be a codesign modeling language (CML) with single semantics, while VHDL is a hardware description language (HDL) with different semantics in simulation and synthesis.

References

[1] G. Agha; "The Structure and Semantics of Actor Languages"; *Lecture Notes in Computer Science, Foundation of Object-Oriented Languages*; Springer-Verlag, 1990.

[2] K. Arnold, J. Gosling; *The Java Programming Language*; Addison-Wesley, 1996.

[3] D. Drusinsky and D. Harel. "Using Statecharts for hardware description and synthesis". In *IEEE Transactions on Computer Aided Design*, 1989.

[4] R. French, M. Lam, J. Levitt, K. Olukotun "A General Method for Compiling Event-Driven Simulation"; *Proceedings of 32th Design Automation Conference*, 6, 1995.

[5] D.D. Gajski, F. Vahid, and S. Narayan. "SpecCharts: a VHDL front-end for embedded systems". UC Irvine, Dept. of ICS, Technical Report 93-31, 1993.

[6] D. Gajski, F. Vahid, S. Narayan, J. Gong. *Specification and Design of Embedded Systems*. New Jersey, Prentice Hall, 1994.

[7] D. Harel; "StateCharts: a Visual Formalism for Complex Systems"; *Science of Programming*, 8, 1987.

[8] E.A. Lee, A. Sangiovanni-Vincentelli; "Comparing Models of Computation"; *Proc. of ICCAD*; San Jose, CA, Nov. 10-14, 1996.

[9] *Common Object Request Broker: Architecture and Specification*; http:// www.omg.org/corbask.htm.

[10] B. Stroustrup; *The C++ Programming Language*; Addison-Wesley, Reading, 1987.