# **15. Scheduling Semantics**

## 15.1 Execution of a hardware model and its verification environment

NOTE—This clause gives an overview of the interactions between these elements, especially with respect to the scheduling and execution of events. Although SystemVerilog is not limited to simulation, the semantics of the language are defined for event directed simulation, and other uses of the hardware description language are abstracted from this base definition.

#### 15.2 Event simulation

The SystemVerilog language is defined in terms of a discrete event execution model. The discrete event simulation is described in more detail in this clause to provide a context to describe the meaning and valid interpretation of SystemVerilog constructs. These resulting definitions provide the standard SystemVerilog reference algorithm for simulation, which all compliant simulators shall implement. Note that there is a great deal of choice in the definitions that follow, and differences in some details of execution are to be expected between different simulators. In addition, SystemVerilog simulators are free to use different algorithms than those described in this clause, provided the user-visible effect is consistent with the reference algorithm.

A SystemVerilog description consists of connected threads of execution or processes. Processes are objects that can be evaluated, that can have state, and that can respond to changes on their inputs to produce outputs. Processes are concurrently scheduled elements, such as initial blocks. Example of processes include, but are not limited to, primitives, initial, always, always\_comb, always\_latch, and always\_ff procedural blocks, continuous assignments, asynchronous tasks, and procedural assignment statements.

Every change in state of a net or variable in the system description being simulated is considered an update event.

Processes are sensitive to update events. When an update event is executed, all the processes that are sensitive to that event are considered for evaluation in an arbitrary order. The evaluation of a process is also an event, known as an evaluation event.

Evaluation events also include PLI callbacks, which are points in the execution model where user-defined external routines can be called from the simulation kernel.

In addition to events, another key aspect of a simulator is time. The term simulation time is used to refer to the time value maintained by the simulator to model the actual time it would take for the system description being simulated. The term time is used interchangeably with simulation time in this clause.

To fully support clear and predictable interactions, a single time slot is divided into multiple regions where events can be scheduled that provide for an ordering of particular types of execution. This allows properties and checkers to sample data when the design under test is in a stable state. Property expressions can be safely evaluated, and testbenches can react to both properties and checkers with zero delay, all in a predictable manner. This same mechanism also allows for non-zero delays in the design, clock propagation, and/or stimulus and response code to be mixed freely and consistently with cycle accurate descriptions.

### 15.3 The stratified event scheduler

A compliant SystemVerilog simulator must maintain some form of data structure that allows events to be dynamically scheduled, executed and removed as the simulator advances through time. The data structure is normally implemented as a time ordered set of linked lists, which are divided and subdivided in a well defined manner.

The first division is by time. Every event has one and only one simulation execution time, which at any given point during simulation can be the current time or some future time. All scheduled events at a specific time define a time slot. Simulation proceeds by executing and removing all events in the current simulation time slot before moving on to the next non-empty time slot, in time order. This procedure guarantees that the simulator never goes backwards in time.

A time slot is divided into a set of ordered regions:

- 1) Preponed
- 2) Pre-active
- 3) Active
- 4) Inactive
- 5) Pre-NBA
- 6) NBA
- 7) Post-NBA
- 8) Observed
- 9) Post-observed
- 10) Reactive
- 11) Re-inactive
- 12) Pre-postponed
- 13) Postponed

The purpose of dividing a time slot into these ordered regions is to provide predictable interactions between the design and testbench code.

Except for the Observed, Reactive, and Re-Inactive regions and the Post-observed PLI region, these regions essentially encompass the Verilog standard reference model for simulation, with exactly the same level of determinism. This means that legacy Verilog code shall continue to run correctly without modification within the new mechanism. The Postponed region is where the monitoring of signals, and other similar events, takes place. No new value changes are allowed to happen in the time slot once the Postponed region is reached.

The Observed, Reactive, and Re-inactive regions are new in the SystemVerilog standard, and events are only scheduled into these new regions from new language constructs.

The Observed region is for the evaluation of the property expressions when they are triggered. A criterion for this determinism is that the property evaluations must only occur once in any clock triggering time slot. During the property evaluation, pass/fail code shall be scheduled in the Reactive region of the current time slot.

The sampling time of sampled data for property expressions is controlled in the clocking block. The new #1step sampling delay provides the ability to sample data immediately before entering the current time slot, and is a preferred construct over other equivalent constructs because it allows the 1step time delay to be parameterized. This #1step construct is a conceptual mechanism that provides a method for defining when sampling takes place, and does not require that an event be created in this previous time slot. Conceptually this #1step sampling is identical to taking the data samples in the Preponed region of the current time slot.

Code specified in the program block, and pass/fail code from property expressions, are scheduled in the Reactive region. A #0 control delay specified in a program block schedules the process for resumption in the Re-inactive region. The Re-inactive region is the program block dual of the Inactive region (see below).

The Pre-active, Pre-NBA, and Post-NBA are new in the SystemVerilog standard but support existing PLI callbacks. The Post-observed region is new in the SystemVerilog standard and has been added for PLI support.

The Pre-active region is specifically for a PLI callback control point that allows for user code to read and write values and create events before events in the Active region are evaluated (see 15.4).

The Pre-NBA region is specifically for a PLI callback control point that allows for user code to read and write values and create events before the events in the NBA region are evaluated (see 15.4).

The Post-NBA region is specifically for a PLI callback control point that allows for user code to read and write values and create events after the events in the NBA region are evaluated (see 15.4).

The Post-observed region is specifically for a PLI callback control point that allows for user code to read values after properties are evaluated (in Observed or earlier region).

The Pre-postponed region is specifically for a PLI callback control point that allows for user code to read and write values and create events after processing all other regions except the Postponed region.

The flow of execution of the event regions is specified in Figure 15-1.



Figure 15-1 — The SystemVerilog flow of time slots and event regions

The Active, Inactive, Pre-NBA, NBA, Post-NBA, Observed, Post-observed, Reactive, Re-inactive, and Prepostponed regions are known as the *iterative* regions.

The Preponed region provides a PLI callback control point that allows user code to access data at the current time slot before any net or variable has changed state. Within this region, it is illegal to write values to any net or variable, or to schedule an event in any other region within the current time slot.

The Active region holds current events being evaluated and can be processed in any order.

The Inactive region holds the events to be evaluated after all the active events are processed.

An explicit #0 delay control suspends the process and schedules an event into the Inactive region (or Re-inactive for program blocks) of the current time slot so that the process is resumed in the next inactive to active iteration.

A nonblocking assignment creates an event in the NBA region, scheduled for current or a later simulation time.

The Postponed region provides a PLI callback control point that allows user code to be suspended until after all the Active, Inactive, NBA, Observe, Reactive, and Re-inactive regions have completed. Within this region, it is illegal to write values to any net or variable, or to schedule an event in any region within the current time slot.

#### 15.3.1 The SystemVerilog simulation reference algorithm

```
execute simulation {
   T = 0;
   initialize the values of all nets and variables;
   schedule all initialization events into time 0 slot;
   while (some time slot is non-empty) {
      move to the next future non-empty time slot and set T;
      execute time slot (T);
   }
}
execute time slot {
   execute region (preponed);
   execute region (pre-active);
   while (any iterative region is non-empty) { /* active ... pre-postponed */
      execute region (active);
      R = first non-empty region in [active ... post-observed];
      if (R is non-empty) {
          move events in R to the active region;
          continue while loop;
      }
      while (reactive or re-inactive region are non-empty)
          execute region(reactive);
          R = first non-empty region in [reactive ... re-inactive];
          if (R is non-empty) {
            move events in R to the reactive region;
          }
       }
       execute region (pre-postponed);
   }
   execute region (postponed);
}
execute region {
   while (region is non-empty) {
      E = any event from region;
      remove E from the region;
      if (E is an update event) {
         update the modified object;
         evaluate processes sensitive to the object and possibly schedule
            further events for execution;
      } else { /* E is an evaluation event */
         evaluate the process associated with the event and possibly
          schedule further events for execution;
      }
   }
}
```

The Iterative regions and their order are: Active, Inactive, Pre-NBA, NBA, Post-NBA, Observed, Post-observed, Reactive, Re-inactive, and Pre-postponed. As shown in the algorithm, once the Reactive or Re-Inactive regions are processed, iteration over the other regions does not resume until these two regions are empty.

## 15.4 The PLI callback control points

There are two kinds of PLI callbacks, those that are executed immediately when some specific activity occurs, and those that are explicitly registered as a one-shot evaluation event.

The following list provides the mapping from the various PLI callbacks

Callback	<b>Event Region</b>
cbAfterDelay	Pre-active
cbNextSimTime	Pre-active
cbReadWriteSynch	Pre-NBA or Post-NBA
cbAtStartOfSimTime	Pre-active
cbNBASynch	Pre-NBA
cbAtEndOfSimTime	Pre-postponed
cbReadOnlySynch	Postponed

#### Table 15-3: PLI Callbacks