

---

As mentioned in the previous chapter, one of the mechanisms that can be used to improve the efficiency and reliability of a process is automation. This chapter covers technologies used in a state-of-the-art functional verification environment. Some of these technologies, such as simulators, are essential for the functional verification activity to take place. Others, such as linting or code coverage technologies, automate some of the most tedious tasks of verification and help increase the confidence in the outcome of the functional verification. This chapter does not contain an exhaustive list of verification technologies, as new application-specific and general purpose verification automation technologies are regularly brought to market.

It is not necessary to use all the technologies mentioned.

As a verification engineer, your job is to use the necessary technologies to ensure that the verification process does not miss a significant functional bug. As a project manager responsible for the delivery of a working product on schedule and within the allocated budget, your responsibility is to arm your engineers with the proper tools to do their job efficiently and with the necessary degree of confidence. Your job is also to decide when the cost of finding the next functional bug has increased above the value the additional functional correctness brings. This last responsibility is the heaviest of them all. Some of these technologies provide information to help you decide when you've reached that point.

A tool may include more than one technology.

This chapter presents various verification technologies separately from each other. Each technology is used in multiple EDA tools. A specific EDA tool may include more than one technology. For example, “super linting” tools leverage linting and formal technologies. Hybrid- or semi-formal tools use a combination of simulation and formal technologies.

Synopsys tools are mentioned.

Being a Synopsys employee at the time of writing this book, the commercial tools I mention are provided by Synopsys, Inc. Other EDA companies supply competitive products. All trademarks and service marks, registered or not, are the property of their respective owners.

---

## LINTING

---

Linting technology finds common programmer mistakes.

The term “lint” comes from the name of a UNIX utility that parses a C program and reports questionable uses and potential problems. When the C programming language was created by Dennis Ritchie, it did not include many of the safeguards that have evolved in later versions of the language, like ANSI-C or C++, or other strongly-typed languages such as Pascal or ADA. *lint* evolved as a tool to identify common mistakes programmers made, letting them find the mistakes quickly and efficiently, instead of waiting to find them through a dreaded segmentation fault during execution of the program.

*lint* identifies real problems, such as mismatched types between arguments and function calls or mismatched number of arguments, as shown in Sample 2-1. The source code is syntactically correct and compiles without a single error or warning using gcc version 2.96.

However, Sample 2-1 suffers from several pathologically severe problems:

1. The *my\_func* function is called with only one argument instead of two.
2. The *my\_func* function is called with an integer value as a first argument instead of a pointer to an integer.

Problems are found faster than at runtime.

As shown in Sample 2-2, the *lint* program identifies these problems, letting the programmer fix them before executing the pro-

---

**Sample 2-1.**  
Syntactically  
correct K&R  
C source code

```
int my_func(addr_ptr, ratio)
    int    *addr_ptr;
    float  ratio;
{
    return (*addr_ptr)++;
}

main()
{
    int my_addr;
    my_func(my_addr);
}
```

gram and observing a catastrophic failure. Diagnosing the problems at runtime would require a debugger and would take several minutes. Compared to the few seconds it took using *lint*, it is easy to see that the latter method is more efficient.

---

**Sample 2-2.**  
*Lint* output for  
Sample 2-1

```
src.c(3): warning: argument ratio unused in
function my_func
src.c(11): warning: addr may be used before set
src.c(12): warning: main() returns random value
to invocation environment
my_func: variable # of args.      src.c(4)  ::
src.c(11)
my_func, arg. 1 used inconsistently
src.c(4)  ::  src.c(11)
my_func returns value which is always ignored
```

Linting does not  
require stimulus.

Linting has a tremendous advantage over other verification technologies: It does not require stimulus, nor does it require a description of the expected output. It performs checks that are entirely static, with the expectations built into the linting tool itself.

## The Limitations of Linting Technology

Linting can only  
identify a certain  
class of prob-  
lems.

Other potential problems were also identified by *lint*. All were fixed in Sample 2-3, but *lint* continues to report a problem with the invocation of the *my\_func* function: The return value is always ignored. Linting cannot identify all problems in source code. It can only find problems that can be statically deduced by looking at the code structure, not problems in the algorithm or data flow.

For example, in Sample 2-3, linting does not recognize that the uninitialized *my\_addr* variable will be incremented in the *my\_func* function, producing random results. Linting is similar to spell checking; it identifies misspelled words, but cannot determine if the wrong word is used. For example, this book could have several instances of the word “with” being used instead of “width”. It is a type of error the spell checker (or a linting tool) could not find.

---

**Sample 2-3.**  
Functionally  
correct K&R  
C source code

```
int my_func(addr_ptr)
    int *addr_ptr;
{
    return (*addr_ptr)++;
}

main()
{
    int my_addr;
    my_func(&my_addr);
    return 0;
}
```

Many false negatives are reported.

Another limitation of linting technology is that it is often too paranoid in reporting problems it identifies. To avoid letting an error go unreported, linting errs on the side of caution and reports potential problems where none exist. This results into a lot of false errors. Designers can become frustrated while looking for non-existent problems and may abandon using linting altogether.

Carefully filter error messages!

You should filter the output of linting tools to eliminate warnings or errors known to be false. Filtering error messages helps reduce the frustration of looking for non-existent problems. More importantly, it reduces the output clutter, reducing the probability that the report of a real problem goes unnoticed among dozens of false reports. Similarly, errors known to be true positive should be highlighted. *Extreme* caution must be exercised when writing such a filter: You must make sure that a true problem is not filtered out and never reported.

Naming conventions can help output filtering.

A properly defined naming convention is a useful technique to help determine if a warning is significant. For example, the report in Sample 2-4 about a latch being inferred on a signal whose name

ends with “\_lat” could be considered as expected and a false warning. All other instances would be flagged as true errors.

---

**Sample 2-4.**  
Output from a  
hypothetical  
SystemVerilog  
linting  
tool

```
Warning: file decoder.sv, line 23: Latch
inferred on reg "address_lat".
Warning: file decoder.sv, line 36: Latch
inferred on reg "next_state".
```

Do not turn off  
checks.

Filtering the output of a linting tool is preferable to turning off checks from within the source code itself or via a command line option. A check may remain turned off for an unexpected duration, potentially hiding real problems. Checks that were thought to be irrelevant may become critical as new source files are added.

Lint code as it is  
being written.

Because it is better to fix problems when they are created, you should run *lint* on the source code while it is being written. If you wait until a large amount of code is written before linting it, the large number of reports—many of them false—will be daunting and create the impression of a setback. The best time to identify a report as true or false is when you are still intimately familiar with the code.

Enforce coding  
guidelines.

Linting, through the use of user-defined rules, can also be used to enforce coding guidelines and naming conventions<sup>1</sup>. Therefore, it should be an integral part of the authoring process to make sure your code meets the standards of readability and maintainability demanded by your audience.

## Linting SystemVerilog Source Code

Linting System-  
Verilog source  
code catches  
common errors.

Linting SystemVerilog source code ensures that all data is properly handled without accidentally dropping or adding to it. The code in Sample 2-5 shows a SystemVerilog model that looks perfect, compiles without errors, but eventually produces incorrect results.

Problems may  
not be obvious.

The problem lies with the use of the *byte* type. It is a signed 8-bit value. It will thus never be equal to or greater than 255 as specified in the conditional expressions. The counter will never saturate and

---

1. See Appendix A for a set of coding guidelines.

---

**Sample 2-5.**  
Potentially  
problematic  
SystemVer-  
ilog code

```
module saturated_counter(output done,
                        input  rst,
                        input  clk);

    byte counter;
    always_ff @(posedge clk)
    begin
        if (rst) counter <= 0;
        else if (counter < 255) counter++;
    end

    assign done = (counter == 255);

endmodule
```

roll over instead. A simple change to “*bit [7:0]*” fixes the problem. But identifying the root cause may be difficult using simulation and waveforms. It is even possible that the problem will never be exercised because no testbench causes the counter to (normally) saturate or the correct effect of the saturation is never checked in the self-checking structure. Linting should report this discrepancy immediately and the bug would be fixed in a few minutes, without a single simulation cycle.

Linting may  
detect some race  
conditions.

Sample 2-6 shows another example. It is a race condition between two concurrent execution branches that will yield an unpredictable result (this race condition is explained in details in the section titled “Write/Write Race Conditions” on page 180). This type of error could be easily detected through linting.

---

**Sample 2-6.**  
Race condition  
in SystemVer-  
ilog code

```
begin
    integer i;
    ...
    fork
        i = 1;
        i = 0;
    join
    ...
end
```

Linting may be  
built in simula-  
tors.

SystemVerilog simulators may provide linting functionality. Some errors, such as race conditions, may be easier to identify during a simulation than through static analysis of the source code. The race condition in Sample 2-6 is quickly identified when using the *+race* command line option of VCS.

Linting tools may leverage formal technology.

Linting tools may use formal technologies to perform more complex static checks. Conversely, property checking tools may also provide some *lint*-like functionality. Some errors, such as unreachable lines of code or FSM transitions, require formal-like analysis of the conditions required to reach executable statements or states and whether or not these conditions can be produced.

## Code Reviews

Reviews are performed by peers.

Although not technically linting, the objective of code reviews is essentially the same: Identify functional and coding style errors before functional verification and simulation. Linting can only identify questionable language uses. It cannot check if the intended behavior has been coded. In code reviews, the source code produced by a designer is reviewed by one or more peers. The goal is not to publicly ridicule the author, but to identify problems with the original code that could not be found by an automated tool. Reviews can identify discrepancies between the design intent and the implementation. They also provide an opportunity for suggesting coding improvements, such as better comments, better structure or the use of assertions.

Identify qualitative problems and functional errors.

A code review is an excellent venue for evaluating the maintainability of a source file, and the relevance of its comments and assertions. Other qualitative coding style issues can also be identified. If the code is well understood, it is often possible to identify functional errors or omissions.

Code reviews are not new ideas either. They have been used for many years in the software design industry. They have been shown to be the most effective quality-enhancing activity. Detailed information on how to conduct effective code reviews can be found in the *resources* section at:

<http://janick.bergeron.com/wtb>

## SIMULATION

---

Simulate your design before implementing it.

Simulation is the most common and familiar verification technology. It is called “simulation” because it is limited to approximating reality. A simulation is never the final goal of a project. The goal of all hardware design projects is to create real physical designs that

	can be sold and generate profits. Simulation attempts to create an artificial universe that mimics the future real design. This type of verification technology lets the designers interact with the design before it is manufactured, and correct flaws and problems earlier.
Simulation is only an approximation of reality.	You must never forget that a simulation is an approximation of reality. Many physical characteristics are simplified—or even ignored—to ease the simulation task. For example, a four-state digital simulation assumes that the only possible values for a signal are 0, 1, unknown, and high-impedance. However, in the physical—and analog—world, the value of a signal is a continuous function of the voltage and current across a thin aluminium or copper wire track: an infinite number of possible values. In a discrete simulation, events that happen deterministically 5 nanoseconds apart may be asynchronous in the real world and may occur randomly.
Simulation is at the mercy of the descriptions being simulated.	Within that simplified universe, the only thing a simulator does is execute a description of the design. The description is limited to a well-defined language with precise semantics. If that description does not accurately reflect the reality it is trying to model, there is no way for you to know that you are simulating something that is different from the design that will be ultimately manufactured. Functional correctness and accuracy of models is a big problem as errors cannot be proven <i>not</i> to exist.

### Stimulus and Response

Simulation requires stimulus.	Simulation is not a static technology. A static verification technology performs its task on a design without any additional information or action required by the user. For example, linting and property checking are static technologies. Simulation, on the other hand, requires that you provide a facsimile of the environment in which the design will find itself. This facsimile is called a testbench. Writing this testbench is the main objective of this book. The testbench needs to provide a representation of the inputs observed by the design, so the simulation can emulate the design's responses based on its description.
The simulation outputs are validated externally, against design intents.	The other thing that you must not forget is that a simulation has no knowledge of your intentions. It cannot determine if a design being simulated is correct. Correctness is a value judgment on the outcome of a simulation that must be made by you, the engineer. Once the design is subjected to an approximation of the inputs from its



environment, your primary responsibility is to examine the outputs produced by the simulation of the design's description and determine if that response is appropriate.

## Event-Driven Simulation


Simulators are never fast enough.

Simulators are continuously faced with one intractable problem: They are never fast enough. They are attempting to emulate a physical world where electricity travels at the speed of light and millions of transistors switch over one billion times in a second. Simulators are implemented using general purpose computers that can execute, under ideal conditions, in the order of a billion sequential instructions per second. The speed advantage is unfairly and forever tipped in favor of the physical world.

Outputs change only when an input changes.

One way to optimize the performance of a simulator is to avoid simulating something that does not need to be simulated. Figure 2-1 shows a 2-input XOR gate. In the physical world, even if the inputs do not change (Figure 2-1(a)), voltage is constantly applied to the output, current is continuously flowing through the transistors (in some technologies), and the atomic particles in the semiconductor are constantly moving. The *interpretation* of the electrical state of the output as a binary value (either a logic 1 or a logic 0) does not change. Only if one of the inputs change (as in Figure 2-1(b)), can the output change.

**Figure 2-1.**  
Behavior of an  
XOR gate



Change in values, called *events*, drive the simulation process.

Sample 2-7 shows a SystemVerilog description (or model) of an XOR gate. The simulator could choose to execute this model continuously, producing the same output value if the input values did not change. An opportunity to improve upon that simulator's performance becomes obvious: do not execute the model while the inputs are constants. Phrased another way: Only execute a model when an input changes. The simulation is therefore driven by

changes in inputs. If you define an input change as an *event*, you now have an *event-driven* simulator.


**Sample 2-7.**  
SystemVerilog model for  
an XOR gate

```
assign Z = A ^ B;
```

Sometimes,  
input changes do  
not cause the  
output to  
change.

But what if both inputs change, as in Figure 2-1(c)? In the logical world, the output does not change. What should an event-driven simulator do? For two reasons, the simulator should execute the description of the XOR gate. First, in the real world, the output of the XOR gate *does* change. The output might oscillate between 0 and 1 or remain in the “neither-0-nor-1” region for a few hundredths of picoseconds (see Figure 2-2). It just depends on how accurate you want your model to be. You could decide to model the XOR gate to include the small amount of time spent in the unknown (or x) state to more accurately reflect what happens when both inputs change at the same time.


**Figure 2-2.**  
Behavior of an  
XOR gate  
when both  
inputs change



Descriptions  
between inputs  
and outputs are  
arbitrary.

The second reason is that the event-driven simulator does not know apriori that it is about to execute a model of an XOR gate. All the simulator knows is that it is about to execute a description of a 2-input, 1-output function. Figure 2-3 shows the view of the XOR gate from the simulator’s perspective: a simple 2-input, 1-output black box. The black box could just as easily contain a 2-input AND gate (in which case the output might very well change if both inputs change), or a 1024-bit linear feedback shift register (LFSR).

**Figure 2-3.**  
Event-driven  
simulator view  
of an XOR  
gate



Acceleration options are often available in event-driven simulators


The mechanism of event-driven simulation introduces some limitations and interesting side effects that are discussed further in Chapter 4.

Simulation vendors are forever locked in a constant battle of beating the competition with an easier-to-use, faster simulator. It is possible to increase the performance of an event-driven simulator by simplifying some underlying assumptions in the design or in the simulation algorithm. For example, reducing delay values to identical unit delays or using two states (0 and 1) instead of four states (0, 1, x and z) are techniques used to speed-up simulation. You should refer to the documentation of your simulator to see what acceleration options are provided. It is also important to understand what the consequences are, in terms of reduced accuracy, of using these acceleration options.

## Cycle-Based Simulation

Figure 2-4 shows the event-driven view of a synchronous circuit composed of a chain of three 2-input gates between two edge-triggered flip-flops. Assuming that Q1 holds a zero, Q2 holds a one and all other inputs remain constant, a rising edge on the clock input would cause an event-driven simulator to simulate the circuit as follows:

**Figure 2-4.**  
Event-driven simulator view of a synchronous circuit



1. The event (rising edge) on the clock input causes the execution of the description of the flip-flop models, changing the output value of Q1 to one and of Q2 to zero, after some small delay.

2. The event on Q1 causes the description of the AND gate to execute, changing the output S1 to one, after some small delay.
3. The event on S1 causes the description of the OR gate to execute, changing the output S2 to one, after some small delay.
4. The event on S2 causes the description of the XOR gate to execute, changing the output S3 to one after some small delay.
5. The next rising edge on the clock causes the description of the flip-flops to execute, Q1 remains unchanged, and Q2 changes back to one, after some small delay.

Many intermediate events in synchronous circuits are not functionally relevant.

To simulate the effect of a single clock cycle on this simple circuit required the generation of six events and the execution of seven models (some models were executed twice). If all we are interested in are the final states of Q1 and Q2, not of the intermediate combinatorial signals, then the simulation of this circuit could be optimized by acting only on the significant events for Q1 and Q2: the active edge of the clock. Phrased another way: Simulation is based on clock cycles. This is how cycle-based simulators operate.

The synchronous circuit in Figure 2-4 can be simulated in a cycle-based simulator using the following sequence:

Cycle-based simulators collapse combinatorial logic into equations.

1. When the circuit description is compiled, all combinatorial functions are collapsed into a single expression that can be used to determine all flip-flop input values based on the current state of the fan-in flip-flops.

For example, the combinatorial function between Q1 and Q2 would be compiled from the following initial description:

$$\begin{array}{rcl} S1 & = & Q1 \ \& \ '1' \\ S2 & = & S1 \ \mid \ '0' \\ S3 & = & S2 \ \wedge \ '0' \end{array}$$


into this final single expression:

$$S3 = Q1$$

The cycle-based simulation view of the compiled circuit is shown in Figure 2-5.

2. During simulation, whenever the clock input rises, the value of all flip-flops are updated using the input value returned by the pre-compiled combinatorial input functions.

**Figure 2-5.**  
Cycle-based  
simulator view  
of a  
synchronous  
circuit



The simulation of the same circuit, using a cycle-based simulator, required the generation of two events and the execution of a single model. The number of logic computations performed is the same in both cases. They would have been performed whether the “A” input changed or not. As long as the time required to perform logic computation is smaller than the time required to schedule intermediate events,<sup>1</sup> and there are many registers changing state at every clock cycle, cycle-based simulation will offer greater performance.

Cycle-based simulations have no timing information.

This great improvement in simulation performance comes at a cost: All timing and delay information is lost. Cycle-based simulators assume that the entire simulation model of the design meets the setup and hold requirements of all the flip-flops. When using a cycle-based simulator, timing is usually verified using a static timing analyzer.

Cycle-based simulators can only handle synchronous circuits.

Cycle-based simulators further assume that the active clock edge is the only significant event in changing the state of the design. All other inputs are assumed to be perfectly synchronous with the active clock edge. Therefore, cycle-based simulators can only simulate perfectly synchronous designs. Anything containing asynchronous inputs, latches or multiple-clock domains *cannot* be simulated accurately. Fortunately, the same restrictions apply to static timing analysis. Thus, circuits that are suitable for cycle-based simulation to verify the functionality are suitable for static timing verification to verify the timing.

## Co-Simulators

No real-world design and testbench is perfectly suited for a single simulator, simulation algorithm or modeling language. Different


1. And it is. By a long shot.

components in a design may be specified using different languages. A design could contain small sections that cannot be simulated using a cycle-based algorithm. Some portion of the design may contain some legacy blocks coded in VHDL or be implemented using analog circuitry.

Multiple simulators can handle separate portions of a simulation.


To handle the portions of a design that do not meet the requirements for cycle-based simulation, most cycle-based simulators are integrated with an event-driven simulator. As shown in Figure 2-6, the synchronous portion of the design is simulated using the cycle-based algorithm, while the remainder of the design is simulated using a conventional event-driven simulator. Both simulators (event-driven and cycle-based) are running together, cooperating to simulate the entire design.

**Figure 2-6.**  
Event-driven  
and cycle-  
based co-  
simulation



Other popular co-simulation environments provide VHDL, SystemVerilog, SystemC, assertions and analog co-simulation. For example, Figure 2-7 shows the testbench (written in SystemVerilog) and a mixed-signal design co-simulated using a SystemVerilog digital simulator and an analog simulator.

**Figure 2-7.**  
Digital and  
analog co-  
simulation



All simulators operate in locked-step.

During co-simulation, all simulators involved progress along the time axis in lock-step. All are at simulation time  $T_1$  at the same time and reach the next time  $T_2$  at the same time. This implies that the speed of a co-simulation environment is limited by the slowest simulator. Some experimental co-simulation environments implement *time warp* synchronization where some simulators are allowed to move ahead of the others.

Performance is decreased by the communication and synchronization overhead.

The biggest hurdle of co-simulation comes from the communication overhead between the simulators. Whenever a signal generated within a simulator is required as an input by another, the current value of that signal, as well as the timing information of any change in that value, must be communicated. This communication usually involves a translation of the event from one simulator into an (almost) equivalent event in another simulator. Ambiguities can arise during that translation when each simulation has different semantics. The difference in semantics is usually present: the semantic difference often being the requirement for co-simulation in the first place.


Translating values and events from one simulator to another can create ambiguities.

Examples of translation ambiguities abound. How do you map SystemVerilog's 128 possible states (composed of orthogonal logic values and strengths) into VHDL's nine logic values (where logic values and strengths are combined)? How do you translate a voltage and current value in an analog simulator into a logic value and strength in a digital simulator? How do you translate an x or z value into a 2-state C++ value?


Co-simulation should not be confused with single-kernel simulation.

Co-simulation is when two (or more) simulators are cooperating to simulate a design, each simulating a portion of the design, as shown in Figure 2-8. It should not be confused with simulators able to read and compile models described in different languages. For example, Synopsys' VCS can simulate a design described using a mix of SystemVerilog, VHDL, OpenVera and SystemC. As shown in Figure 2-9, all languages are compiled into a single internal representation or machine code and the simulation is performed using a single simulation engine.

**Figure 2-8.**  
Co-simulator



**Figure 2-9.**  
Mixed-  
language  
simulator



**VERIFICATION INTELLECTUAL PROPERTY**

You can buy IP for standard functions.

If you want to verify your design, it is necessary to have models for all the parts included in a simulation. The model of the RTL design is a natural by-product of the design exercise and the actual objective of the simulation. Models for embedded or external RAMs are also required, as well as models for standard interfaces and off-the-shelf parts. If you were able to procure the RAM, design IP, specification or standard part from a third party, you should be able to obtain a model for it as well. You may have to obtain the model from a different vendor than the one who supplies the design component.

It is cheaper to buy models than write them yourself.

At first glance, buying a simulation model from a third-party provider may seem expensive. Many have decided to write their own models to save on licensing costs. However, you have to decide if this endeavor is truly economically fruitful: Are you in the modeling business or in the chip design business? If you have a shortage of qualified engineers, why spend critical resources on writing a model that does not embody any competitive advantage for your company? If it was not worth designing on your own in the first place, why is writing your own model suddenly justified?



Your model is not as reliable as the one you buy.

Secondly, the model you write has never been used before. Its quality is much lower than a model that has been used by several other companies before you. The value of a functionally correct and reliable model is far greater than an uncertain one. Writing and verifying a model to the *same degree of confidence* as the third-party model is always more expensive than licensing it. And be assured: No matter how simple the model is (such as a quad 2-input NAND gate, 74LS00), you'll get it wrong the first time. If not functionally, then at least with respect to timing or connectivity.


There are several providers of verification IP. Many are written using a proprietary language or C code; others are provided as non-synthesizeable SystemVerilog source code. For intellectual property protection and licensing technicalities, most are provided as compiled binary or encrypted models. Verification IP includes, but is not limited to functional models of external and embedded memories, bus-functional models for standard interfaces, protocol generators and analyzers, assertion sets for standard protocols and black-box models for off-the-shelf components and processors.

## WAVEFORM VIEWERS

Waveform viewers display the changes in signal values over time.

Waveform viewing is the most common verification technology used in conjunction with simulators. It lets you visualize the transitions of multiple signals over time, and their relationship with other transitions. With a waveform viewer, you can zoom in and out over particular time sequences, measure time differences between two transitions, or display a collection of bits as bit strings, hexadecimal or as symbolic values. Figure 2-10 shows a typical display of a waveform viewer showing the inputs and outputs of a 4-bit synchronous counter.

**Figure 2-10.** Hypothetical waveform view of a 4-bit synchronous counter



Waveform viewing is used to debug simulations.

Waveform viewing is indispensable during the authoring phase of a design or a testbench. With a viewer, you can casually inspect that the behavior of the code is as expected. They are needed to diagnose, in an efficient fashion, why and when problems occur in the design or testbench. They can be used interactively during the simulation, but more importantly offline, after the simulation has completed. As shown in Figure 2-11, a waveform viewer can play back the events that occurred during the simulation that were recorded in some trace file.

Recording waveform trace data decreases simulation performance.

Viewing waveforms as a post-processing step lets you quickly browse through a simulation that can take hours to run. However, keep in mind that recording trace information significantly reduces the performance of the simulator. The quantity and scope of the signals whose transitions are traced, as well as the duration of the trace, should be limited as much as possible. Of course, you have to trade-off the cost of tracing a greater quantity or scope of signals versus the cost of running the simulation over again to get a trace of additional signals that turn out to be required to completely diagnose the problem. If it is likely or known that bugs will be reported, such as the beginning of the project or during a debugging iteration, trace all the signals required to diagnose the problem. If no errors are expected, such as during regression runs, no signal should be traced.

Do not use waveform viewing to determine if a design

Some viewers can compare sets of waveforms.

Some waveform viewers can compare two sets of waveforms. One set is presumed to be a golden reference, while the other is verified for any discrepancy. The comparator visually flags or highlights any differences found. This approach has two significant problems.

How do you define a set of waveforms as “golden”?

First, how is the golden reference waveform set declared “golden”? If visual inspection is required, the probability of missing a significant functional error remains equal to one hundred percent in most cases. The only time golden waveforms are truly available is in a redesign exercise, where cycle-accurate backward compatibility must be maintained. However, there are very few of these designs. Most redesign exercises take advantage of the process to introduce needed modifications or enhancements, thus tarnishing the status of the golden waveforms.

And are the differences really significant?

Second, waveforms are at the wrong level of abstraction to compare simulation results against design intent. Differences from the golden waveforms may not be significant. The value of all output signals is not significant all the time. Sometimes, what is significant is the relative relationships between the transitions, not their absolute position. The new waveforms may be simply shifted by a few clock cycles compared to the reference waveforms, but remain functionally correct. Yet, the comparator identifies this situation as a mismatch.

Use assertions instead.

You should avoid using waveform viewing to check a response. It should be reserved for debugging. Instead of looking for specific signal relationships in a waveform viewer, specify these same relationships using assertions. The assertions will be continuously and reliably checked, for the entire duration of the simulation, for all simulations. They will provide a specification of the intended functionality of the design. Should your design be picked up by another designer, your original intent will be communicated along with your original code.

---

## CODE COVERAGE

---

Did you forget to verify some function in your code?


Code coverage is a technology that can identify what code has been (and more importantly *not* been) executed in the design under verification. It is a technology that has been in use in software engineering for quite some time. The problem with a design containing an unknown bug is that it looks just like a perfectly good design. It is

impossible to know, with one hundred percent certainty, that the design being verified is indeed functionally correct. All of your testbenches simulate successfully, but are there sections of the RTL code that you did not exercise and therefore not triggered a functional error? That is the question that code coverage can help answer.

Code must first be instrumented.

Figure 2-12 shows how code coverage works. The source code is first *instrumented*. The instrumentation process simply adds checkpoints at strategic locations of the source code to record whether a particular construct has been exercised. The instrumentation mechanism varies from tool to tool. Some may use file I/O features available in the language (i.e., use *\$write* statements in SystemVerilog). Others may use special features built into the simulator.

**Figure 2-12.**  
Code coverage process



No need to instrument the testbenches.

Only the code for the design under verification needs to be covered and thus instrumented. The objective of code coverage is to determine if you have forgotten to exercise some code in the design. The code for the testbenches need not be traced to confirm that it has executed. If a significant section of a testbench was not executed, it should be reflected in some portion of the design not being exercised. Furthermore, a significant portion of testbench code is executed only if errors are detected. Code coverage metrics on testbench code are therefore of little interest.

Trace information is collected at runtime.

The instrumented code is then simulated normally using all available, uninstrumented, testbenches. The cumulative traces from all simulations are collected into a database. From that database, reports can be generated to measure various coverage metrics of the verification suite on the design.

The most popular metrics are statement, path and expression coverage.

## Statement Coverage

Statement and block coverage are the same thing.

Statement coverage can also be called block coverage, where a block is a sequence of statements that are executed if a single statement is executed. The code in Sample 2-8 shows an example of a statement block. The block named *acked* is executed entirely whenever the expression in the *if* statement evaluates to *true*. So counting the execution of that block is equivalent to counting the execution of the four individual statements within that block.

**Sample 2-8.**  
Block vs.  
statement execution

```
if (dtack == 1'b1) begin: acked
    as    <= 1'b0;
    data  <= 16'hZZZZ;
    bus_rq <= 1'b0;
    state <= IDLE;
end: acked
```

But block boundaries may not be that obvious.

Statement blocks may not be necessarily clearly delimited. In Sample 2-9, two statement blocks are found: one before (and including) the *wait* statement, and one after. The *wait* statement may have never completed and the execution was waiting forever. The subsequent sequential statements may not have executed. Thus, they form a separate statement block.

**Sample 2-9.**  
Blocks separated by a *wait* statement

```
address <= 16'hFFED;
ale     <= 1'b1;
rw      <= 1'b1;
wait (dtack == 1'b1);
read_data = data;
ale     <= 1'b0;
```

Did you execute all the statements?

Statement, line or block coverage measures how much of the total lines of code were executed by the verification suite. A graphical user interface usually lets the user browse the source code and quickly identify the statements that were not executed. Figure 2-13 shows, in a graphical fashion, a statement coverage report for a small portion of code from a model of a modem. The actual form of

the report from any code coverage tool or source code browser will likely be different.

**Figure 2-13.**  
Example of  
statement  
coverage

```
☒ if (parity == ODD || parity == EVEN) begin
☐   tx <= compute_parity(data, parity);
☐   #(tx_time);
      end
☒ tx <= 1'b0;
☒ #(tx_time);
☒ if (stop_bits == 2) begin
☒   tx <= 1'b0;
☒   #(tx_time);
      end
```

Why did you not  
execute all state-  
ments?

The example in Figure 2-13 shows that two out of the eight executable statements—or 25 percent—were not executed. To bring the statement coverage metric up to 100 percent, a desirable goal<sup>1</sup>, it is necessary to understand what conditions are required to cause the execution of the uncovered statements. In this case, the parity must be set to either ODD or EVEN. Once the conditions have been determined, you must understand why they never occurred in the first place. Is it a condition that can never occur? Is it a condition that should have been verified by the existing verification suite? Or is it a condition that was forgotten?

Add testcases to  
execute all state-  
ments.

If the conditions that would cause the uncovered statements to be executed should have been verified, it is an indication that one or more testbenches are either not functionally correct or incomplete. If the condition was entirely forgotten, it is necessary to add to an existing testbench, create an entirely new one or make additional runs with different seeds.

## Path Coverage

There is more  
than one way to  
execute a  
sequence of  
statements.


Path coverage measures all possible ways you can execute a sequence of statements. The code in Sample 2-10 has four possible paths: the first *if* statement can be either true or false. So can the second. To verify all paths through this simple code section, it is

---

1. But not necessarily achievable. For example, the *default* clause in a fully specified *case* statement should never be executed.

necessary to execute it with all possible state combinations for both *if* statements: false-false, false-true, true-false, and true-true.

**Sample 2-10.**  
Example of  
statement and  
path coverage



Why were some  
sequences not  
executed?

The current verification suite, although it offers 100 percent statement coverage, only offers 75 percent path coverage through this small code section. Again, it is necessary to determine the conditions that cause the uncovered path to be executed. In this case, a testcase must set the parity to neither ODD nor EVEN and the number of stop bits to two. Again, the important question one must ask is whether this is a condition that will ever happen, or if it is a condition that was overlooked.

Limit the length  
of statement  
sequences.

The number of paths in a sequence of statements grows exponentially with the number of control-flow statements. Code coverage tools give up measuring path coverage if their number is too large in a given code sequence. To avoid this situation, keep all sequential code constructs (*always* and *initial* blocks, *tasks* and *functions*) to under 100 lines.

Reaching 100 percent path coverage is very difficult.

## Expression Coverage

There may be  
more than one  
cause for a con-  
trol-flow  
change.

If you look closely at the code in Sample 2-11, you notice that there are two mutually independent conditions that can cause the first *if* statement to branch the execution into its *then* clause: parity being set to either ODD or EVEN. Expression coverage, as shown in Sample 2-11, measures the various ways decisions through the code

are made. Even if the statement coverage is at 100 percent, the expression coverage is only at 50 percent.

**Sample 2-11.**  
Example of  
statement and  
expression  
coverage

```

✓ if (parity == ODD || parity == EVEN) begin
✓   tx <= compute_parity(data, parity);
✓   #(tx_time);
✓ end
✓ tx <= 1'b0;
✓ #(tx_time);
✓ if (stop_bits == 2) begin
✓   tx <= 1'b0;
✓   #(tx_time);
✓ end
✓

```

Once more, it is necessary to understand why a controlling term of an expression has not been exercised. In this case, no testbench sets the parity to EVEN. Is it a condition that will never occur? Or was it another oversight?

Reaching 100 percent expression coverage is extremely difficult.

## FSM Coverage

Statement coverage detects unvisited states.


Because each state in an FSM is usually explicitly coded using a choice in a *case* statement, any unvisited state will be clearly identifiable through uncovered statements. The state corresponding to an uncovered *case* statement choice was not visited during verification.

FSM coverage identifies state transitions.

Figure 2-14 shows a bubble diagram for an FSM. Although it has only five states, it has significantly more possible transitions: 14 possible transitions exist between adjoining states. State coverage of 100 percent can be easily reached through the sequence Reset, A, B, D, then C. However, this would yield only 36 percent transition coverage. To completely verify the implementation of this FSM, it is necessary to ensure the design operates according to expectation for all transitions.



**Figure 2-14.**  
Example FSM  
bubble  
diagram



FSM coverage cannot identify unintended or missing transitions.

The transitions identified by FSM coverage tools are automatically extracted from the implementation of the FSM. There is no way for the coverage tool to determine whether a transition was part of the intent, or if an intended transition is missing. It is important to review the extracted state transitions to ensure that only and all intended transitions are present.

What about unspecified states?

The FSM illustrated in Figure 2-14 only shows five specified states. Once synthesized into hardware, a 3-bit state register will be necessary (maybe more if a different state encoding scheme, such as one-hot, is used). This leaves three possible state values that were not specified. What if some cosmic rays zap the design into one of these unspecified states? Will the correctness of the design be affected? Logic optimization may yield decode logic that creates an island of transitions within those three unspecified states, never letting the design recover into specified behavior unless reset is applied. The issues of design safety and reliability and techniques for ensuring them are beyond the scope of this book. But it is the role of a verification engineer to ask those questions.

Formal verification may be better suited for FSM verification.


The behavior of a FSM is a combination of its state transition description and the behavior of its input signals. If those input signals are themselves generated by another FSM or follow a specific protocol, it is possible that certain transitions cannot be taken or states cannot be reached. A property checker tool may be able to formally determine which states are reachable and which transitions are possible—including invalid states and transitions. It may also be able to formally verify that a specific state encoding, such as *one-hot*, is never violated.

## What Does 100 Percent Code Coverage Mean?

Completeness  
does not imply  
correctness.

The short answer is: The entire design implementation was executed. Code coverage indicates how *thoroughly* your entire verification suite exercises the source code. But it does not provide an indication, in any way, about the *correctness* or *completeness* of the verification suite. Figure 2-15 shows the reconvergence model for automatically extracted code coverage metrics. It clearly shows that it does not help verify design intent, only that the RTL code, correct or not, was fully exercised.

**Figure 2-15.**  
Reconvergent  
paths in  
automated  
code coverage



Results from code coverage should be interpreted with a grain of salt. They should be used to help identify corner cases that were not exercised by the verification suite or implementation-dependent features that were introduced during the implementation. You should also determine if the uncovered cases are relevant and deserve additional attention, or a consequence of the mindlessness of the coverage tool.

Code coverage  
lets you know if  
you are not  
done.

Code coverage indicates if the verification task is *not* complete through low coverage numbers. A high coverage number is by no means an indication that the job is over. For example, the code in an empty *module* will always be 100 percent covered. If the functionality that ought to be implemented in that *module* is not verified, all testbenches will pass without errors. Code coverage is an additional indicator for the completeness of the verification job. It can help increase your confidence that the verification job is complete, but it should not be your only indicator.

Code coverage  
tools can be  
used as profil-  
ers.

When developing models for simulation only, where performance is an important criteria, code coverage can be used for *profiling*. The aim of profiling is the opposite of code coverage. The aim of profiling is to identify the lines of codes that are executed most often. These lines of code become the primary candidates for performance optimization efforts.

---

**FUNCTIONAL COVERAGE**

---

Did you forget to verify some condition?

Functional coverage is another technology to help ensure that a bad design is not hiding behind passing testbenches. Although this technology has been in use at some companies for quite some time, it is a recent addition to the arsenal of general-purpose verification tools. Functional coverage records relevant metrics (e.g., packet length, instruction opcode, buffer occupancy level) to ensure that the verification process has exercised the design through all of the interesting values. Whereas code coverage measures how much of the implementation has been exercised, functional coverage measures how much of the original design specification has been exercised.

It complements code coverage.

High functional coverage does not necessarily correlate with high code coverage. Whereas code coverage is concerned with recording the mechanics of code execution, functional coverage is concerned with the intent or purpose of the implemented function. For example, the decoding of a CPU instruction may involve separate *case* statements for each field in the opcode. Each *case* statement may be 100 percent code-covered due to combinations of field values from previously decoded opcodes. However, the particular combination involved in decoding a specific CPU instruction may not have been exercised.

It will detect errors of omission.

Sample 2-12 shows a *case* statement decoding a CPU instruction. Notice how the decoding of the RTS instruction is missing. If you rely solely on code coverage, you will be lulled in a false sense of completeness by having 100 percent coverage of this code. For code coverage to report a gap, the unexercised code must apriori exist. Functional coverage does not rely on actual code. It will report gaps in the recorded values whether the code to process them is there or not.

---

**Sample 2-12.**  
Example of coding error undetectable by code coverage

```
enum {ADD, SUB, JMP, RTS, NOP} opcode;
...
case (opcode)
  ADD: ...
  SUB: ...
  JMP: ...
  default: ...
endcase
```

It must be manually defined.

Code coverage was quickly adopted into verification processes because of its low adoption cost. It requires very little additional action from the user: usually the specification of an additional command-line option when compiling your code. Functional coverage, because it is a measure of values deemed to be interesting and relevant, must be manually specified. Since relevance and interest are qualities that are extracted from the intent of the design, functional coverage is not something that can be automatically extracted from the RTL source code. Your functional coverage metrics will be only as good as what you implement.

Metrics are collected at runtime and graded.

Like code coverage, functional coverage metrics are collected at runtime, during a simulation run. The values from individual runs are collected into a database or separate files. The functional coverage metrics from these separate runs are then merged for offline analysis. The marginal coverage of individual runs can then be graded to identify which runs contributed the most toward the overall functional coverage goal. These runs are then given preference in the regression suite, while pruning runs that did not significantly contribute to the objective.

Coverage data can be used at runtime.

SystemVerilog provides a set of predefined methods that let a testbench dynamically query a particular functional coverage metric. The testbench can then use the information to modify its current behavior. For example, it could increase the probability of generating values that have not been covered yet. It could decide to abort the simulation should the functional coverage not have significantly increased since the last query.

Although touted as a powerful mechanism, it is no silver bullet. Implementing the dynamic feedback mechanism is not easy: You have to correlate your stimulus generation process with the functional coverage metric, and ensure that one will cause the other to converge toward the goal. Dynamic feedback works best when there is a direct correlation between the input and the measured coverage, such as instruction types. It may be more efficient to achieve your goal with three or four runs of a few simple testbenches without dynamic feedback than with a single run of a much more complex testbench.

---

## Coverage Points

Did I generate all interesting and relevant values?

A coverage point is the sampling of an individual scalar value or expression. The objective of a coverage point is to ensure that all interesting and relevant values have been observed in the sampled value or expression. Examples of coverage points include, but are not limited to, packet length, instruction opcode, interrupt level, bus transaction termination status, buffer occupancy level, bus request patterns and so on.

Define *what* to sample.

It is extremely easy to record functional coverage and be inundated with vast amounts of coverage data. But data is not the same thing as information. You must restrict coverage to only (but all!) values that will indicate how thoroughly your design has been verified. For example, measuring the value of the read and write pointers in a FIFO is fine if you are concerned about the full utilization of the buffer space and wrapping around of the pointer values. But if you are interested in the FIFO occupancy level (Was it ever empty? Was it ever full? Did it overflow?), you should measure and record the *difference* between the pointer values.

Define *where* to sample it.

Next, you must decide where in your testbench or design is the measured value accurate and relevant. For example, you can sample the opcode of an instruction at several places: at the output of the code generator, at the interface of the program memory, in the decoder register or in the execution pipeline. You have to ensure that a value, once recorded, is indeed processed or committed as implied by the coverage metric.

For example, if you are measuring opcodes that were executed, they should be sampled in the execution unit. Sampling them in the decode unit could result in false samples when the decode pipeline is flushed on branches or exceptions. Similarly, sampling the length of packets at the output of the generator may yield false samples: If a packet is corrupted by injecting an error during its transmission to the design in lower-level functions of the testbench, it may be dropped.

Define *when* to sample it.

Values are sampled at some point in time during the simulation. It could be at every clock cycle, whenever the address strobe signal is asserted, every time a request is made or after randomly generating a new value. You must carefully choose your sampling time. Over-

sampling will decrease simulation performance and consume database resources without contributing additional information.

The sampled data must also be stable so race conditions must be avoided between the sampled data and the sampling event (see “Read/Write Race Conditions” on page 177). To reduce the probability that a transient value is being sampled, functional coverage in SystemVerilog can be sampled at the end of the simulation cycle, before time is about to advance (see “The Simulation Cycle” on page 163) when the *strobe* option is used.

Define *why* it is covered.

If functional coverage is supposed to measure interesting and relevant values, it is necessary to define what makes those values so interesting and relevant. For example, measuring the functional coverage of a 32-bit address will yield over 4 billion “interesting and relevant” values. Not all values are created equal—but most are. Values may be numerically different but functionally equivalent. By identifying those functionally equivalent values into a single bin, you can reduce the number of interesting and relevant values to a more manageable size. For example, based on the decoder architecture, addresses 0x00000001 through 0x7FFFFFFF and addresses 0x80000000 through 0x8FFFFFFE are functionally equivalent, reducing the number of relevant and interesting values to 4 bins (min, 1 to mid, mid to max-1, max).

It can detect invalid values.

Just as you can define bins of equivalent values, it is possible to define bins of invalid or unexpected values. Functional coverage can be used as an error detecting mechanism, just like an *if* statement in your testbench code. However, you should not rely on functional coverage to detect invalid values. Functional coverage is an optional runtime data collection mechanism that may not be turned on at all times. If functional coverage is not enabled to improve simulation performance and if a value is defined as invalid in the functional coverage only, then an invalid value may go undetected.

It can report holes.

The ultimate purpose of functional coverage is to identify what remains to be done. During analysis, the functional coverage reporting tool can compare the number of bins that contain at least one sample against the total number of bins. Any bin that does not contain at least one sample is a hole in your functional coverage. By enumerating the empty bins, you can focus on the holes in your test cases and complete your verification sooner rather than continue to exercise functionality that has already been verified.

For this enumeration to be possible, the total number of bins for a coverage point must be relatively small. For example, it is practically impossible to fill the coverage for a 32-bit value without broad bins. The number of holes will be likely in the millions, making enumeration impossible. You should strive to limit the number of possible bins as much as possible.

## Cross Coverage

Did I generate all interesting *combinations* of values?

Whereas coverage points are concerned with individual scalar values, cross coverage measures the presence or occurrence of combinations of values. It helps answer questions like, “Did I inject a corrupted packet on all ports?” “Did we execute all combinations of opcodes and operand modes?” and “Did this state machine visit each state while that buffer was full, not empty and empty?” Cross coverage can involve more than two coverage points. However, the number of possible bins grows factorially with the number of crossed points.

Similar to coverage points.

Mechanically, cross coverage is identical to coverage points. Specific values are sampled at specific locations at specific points in time with specific value bins. The only difference is that two or more values are sampled instead of one. To ensure that crossed values are consistent, they must all be sampled at the same time. In SystemVerilog, only coverage points specified within the same *covergroup* can be crossed.

## Transition Coverage

Did I generate all interesting *sequences* of values?

Whereas cross coverage is concerned with combinations of scalar values at the same point in time, transition coverage measures the presence or occurrence of sequences of values. Transition coverage helps answer questions like, “Did I perform all combinations of back-to-back read and write cycles?” “Did we execute all combinations of arithmetic opcodes followed by test opcodes?” and “Did this state machine traverse all significant paths?” Transition coverage can involve more than two consecutive values of the same coverage point. However, the number of possible bins grows factorially with the number of transition states.

Similar to coverage points.

Mechanically, transition coverage is identical to coverage points. Specific values are sampled at specific locations at specific points

in time with specific bins. The only difference is that a sample is said to have occurred in a bin after two or more consecutive coverage point samples instead of one. The other difference is that transitions can overlap, hence two transition samples may be composed of the same coverage point sample.

Similar to FSM path coverage.

Conceptually, transition coverage is identical to FSM path coverage (see “FSM Coverage” on page 46). Both record the consecutive values at a particular location of the design (for example, a state register), and both compare against the possible set of paths. But unlike FSM coverage tools, which are limited to state registers in RTL code, transition coverage can be applied to any coverage points in testbenches and the design under verification.

Transition coverage reflects intent.

Because transition coverage is manually specified from the intent of the design or the implementation, it provides a true independent path to verifying the correctness of the design and the completeness of the verification. It can detect invalid transitions as well as specify transitions that may be missing from the implementation of the design.

## What Does 100 Percent Functional Coverage Mean?

It indicates completeness of the test suite, not correctness.

Functional coverage indicates which interesting and relevant conditions were verified. It provides an indication of the *thoroughness* of the implementation of the verification plan. Unless some bins are defined as invalid, it cannot provide an indication, in any way, about the *correctness* of those conditions or of the design’s response to those conditions. Functional coverage metrics are only as good as the functional coverage model you have defined. An overall functional coverage metric of 100 percent means that you’ve covered all of the coverage points you included in the simulation. It makes no statement about the *completeness* of your functional coverage model.

Results from functional coverage should also be interpreted with a grain of salt. Since they are generated by additional testbench constructs, they have to be debugged and verified for correctness before being trusted. They will help identify additional interesting conditions that were not included in the verification plan.



Functional coverage lets you know if you are done.

When used properly, functional coverage becomes a formal specification of the verification plan. Once you reach 100 percent functional coverage, it indicates that you have created and exercised all of the relevant and interesting conditions you originally identified. It confirms that you have implemented everything in the verification plan. However, it does not provide any indication of the completeness of the verification plan itself or the correctness of the design under such conditions.

If a metric is not interesting, don't measure it.

It is extremely easy to define functional coverage metrics and generate many reports. If coverage is not measured according to a specific purpose, you will soon drown under megabytes of functional coverage reports. And few of them will ever be close to 100 percent. It will also become impossible to determine which report is significant or what is the significance of the holes in others. The verification plan (see the next chapter) should serve as the functional specification for the coverage models, as well as for the rest of the verification environment. If a report is not interesting or meaningful to look at, if you are not eager to look at a report after a simulation run, then you should question its existence.

---

## VERIFICATION LANGUAGE TECHNOLOGIES

Verilog is a simulation language, not a verification language.

Verilog was designed with a focus on describing low-level hardware structures. Verilog-2001 only introduced support for basic high-level data structures. Verilog thus continued to lack features important in efficiently implementing a modern verification process. These shortcomings were the forces being the creation of hardware verification languages, such as Synopsys' OpenVera. Having demonstrated their usefulness, the value-add functionality of HVLs has been incorporated in SystemVerilog.

Verification languages can raise the level of abstraction.

As mentioned in Chapter 1, one way to increase productivity is to raise the level of abstraction used to perform a task. High-level languages, such as C or Pascal, raised the level of abstraction from assembly-level, enabling software engineers to become more productive. Similarly, the SystemVerilog verification constructs are able to raise the level of abstraction compared to plain Verilog. SystemVerilog can provide an increase in level of abstraction while maintaining the important concepts necessary to interact with hardware: time, concurrency and instantiation. The SystemVerilog fea-


SystemVerilog can automate verification.

tures that help raise the level of abstraction include: *class*, object-oriented *class* extensions and temporal properties.

If higher levels of abstraction and object-orientedness were sufficient, then C++ would have long been identified as the best solution<sup>1</sup>: It is free and widely known. SystemVerilog provides additional benefits, as shown in Figure 2-16. It can automate a portion of the verification process by randomly generating stimulus, collecting functional coverage to identify holes then easily add or modify constraints to create more stimulus targeted to fill those holes. To support this productivity cycle, SystemVerilog offers constrainable random generation, functional coverage measurement and an object-oriented code extension mechanism.

---

**Figure 2-16.** SystemVerilog productivity cycle



SystemVerilog can implement a coverage-driven constrained random approach.

SystemVerilog can be used as if it was a simple souped-up version of Verilog. SystemVerilog will make implementing directed testbenches easier than plain Verilog—especially the self-checking part. But if you want to take advantage of the productivity cycle shown in Figure 2-16, the verification process must be approached—and implemented—in a different fashion.

This change is just like taking advantage of the productivity offered by logic synthesis tools: It requires an approach different from schematic capture. To successfully implement a coverage-driven constrained random verification approach, you need to modify the way you plan your verification, design its strategy and implement the testcases. This new approach is described in “Coverage-Driven Random-Based Approach” on page 101.

---

1. C++ still lacks a native concept of time, concurrency and instantiation.

## ASSERTIONS

---

Assertions detect conditions that should always be true.

An assertion boils down to an *if* statement and an error message should the expression in the *if* statement become false. Assertions have been used in software design for many years: the *assert()* function has been part of the ANSI C standard from the beginning. In software for example, assertions are used to detect conditions such as NULL pointers or empty lists. VHDL has had an *assert* statement from day one too, but it was never a popular construct.

Hardware assertions require a form of temporal language.

An immediate assertion, like an *if* statement, simply checks that, at the time it is executed, the condition evaluates to TRUE. This simple zero-time test is not sufficient for supporting assertions in hardware designs. In hardware, functional correctness usually involves behavior over a period of time. Some hardware assertions such as, “This state register is one-hot encoded.” or “This FIFO never overflows.” can be expressed as immediate, zero-time expressions. But checking simple hardware assertions such as, “This signal must be asserted for a single clock period.” or “A request must always be followed by a grant or abort within 10 clock cycles.” require that the assertion condition be evaluated over time. Thus, assertions require the use of a temporal language to be able to describe relationships over time.

There are two classes of assertions.

Assertions fall in two broad classes: those specified by the designer and those specified by the verification engineer.

- Implementation assertions are specified by the designers.
- Specification assertions are specified by the verification engineers.

Implementation assertions verify assumptions.

*Implementation assertions* are used to formally encode the designer’s assumptions about the interface or implementation of the design or conditions that are indications of misuse or design faults. For example, the designer of a FIFO would add assertions to detect if it ever overflows or underflows or that, because of a design limitation, the *write* and *read* pulses are ever asserted at the same time. Because implementation assertions are specified by the designer, they will not detect discrepancies between the functional intent and the design. But implementation assertions will detect discrepancies between the design assumptions and the implementation.

Specification assertions verify intent.

*Specification assertions* formally encode expectations of the design based on the functional intent. These assertions are used as a functional error detection mechanism and supplement the error detections performed in the self-checking section of testbenches. Specification assertions are typically *white-box* strategies because the relationships between the primary inputs and outputs of a modern design are too complex to be described efficiently in SystemVerilog's temporal languages. For example, rather than relying on the scoreboard to detect that an arbiter is not fair, it is much simpler to perform this check using a white-box assertion.

### Simulated Assertions

The OVL started the storm.

Assertions took the hardware design community by storm when Foster and Bening's book<sup>1</sup> introduced the concept using a library of predefined Verilog modules that implemented a set of common design assertions. The library, available in source form as the *Open Verification Library*,<sup>2</sup> was a clever way of using Verilog to specify temporal expressions. Foster, then at Hewlett-Packard, had a hidden agenda: Get designers to specify design assertions he could then try to prove using formal methods. Using Verilog modules was a convenient solution to ease the adoption of these assertions by the designers. The reality of what happened next proved to be even more fruitful.

They detect errors close in space and time to the fault.

If a design assumption is violated during simulation, the design will not operate correctly. The cause of the violation is not important: It could be a misunderstanding by the designer of the block or the designer of the upstream block or an incorrect testbench. The relevant fact is that the design is failing to operate according to the original intent. The symptoms of that low-level failure are usually not visible (if at all) until the affected data item makes its way to the outputs of the design and is flagged by the self-checking structure.

An assertion formally encoding the design assumption immediately fires and reports a problem at the time it occurs, in the area of the design where it occurs. Debugging and fixing the assertion failure

---

1. Harry Foster and Lionel Bening, "*Principles of Verifiable RTL Design*," second edition, Kluwer Academic Publisher, ISBN 0-7923-7368-5.

2. See <http://www.eda.org/ovl>.

(whatever the cause) will be a lot more efficient than tracing back the cause of a corrupted packet. In one of Foster’s projects, 85% of the design errors were caught and quickly fixed using simulated assertions.

Your model can tell you if things are not as assumed.

SystemVerilog provides a powerful assertion language. But it also provides constructs designed to ensure consistent results between synthesis and simulation. Sample 2-14 shows an example of a synthesizable *unique case* statement, which can be used to replace the *full case* directive shown in Sample 2-13. In both cases, the synthesis tool is instructed that the *case* statement describes all possible non-overlapping conditions. But it is possible for an unexpected condition to occur during simulation. If that were the case, the simulation results would differ from the results produced by the hardware implementation. If a pragma is used, as in Sample 2-13, the unexpected condition would only be detected if it eventually produces an incorrect response. If the *unique case* statement is used, any unexpected condition will be immediately reported near the time and place of its occurrence.

---

**Sample 2-13.**  
*full case* directive

```
case (mode[1:0]) // synopsys full_case
  2'b00: ...
  2'b10: ...
  2'b01: ...
endcase
```

---

**Sample 2-14.**  
*unique case* statement

```
unique case (mode[1:0])
  2'b00: ...
  2'b10: ...
  2'b01: ...
endcase
```

## Formal Assertion Proving

Is it possible for an assertion to fire?

Simulation can show only the presence of bugs, never prove their absence. The fact that an assertion has never reported a violation throughout a series of simulations does not mean that it can never be violated. Tools like code and functional coverage can satisfy us that a portion of a design was thoroughly verified—but there will (and should) always be a nagging doubt.

Property checking can mathematically prove or disprove an assertion.

Formal tools called *property checkers* or *assertion provers* can mathematically prove that, given an RTL design and some assumptions about the relationships of the input signals, an assertion will always hold true. If a counter example is found, the formal tool will provide details on the sequence of events that leads to the assertion violation. It is then up to you to decide if this sequence of events is possible, given additional knowledge about the environment of the design.

Some assertions are used as assumptions.

Given total freedom over the inputs of a design, it may be possible to violate assertions about its implementation. The proper operation of the design may rely on the proper behavior of the inputs, subject to limitations and rules that must be followed. These input signals usually come from other designs that do not behave (one hopes!) erratically and follow the rules. When proving some assertions on a design, it is thus necessary to supply assertions on the inputs or state of the design. The latter assertions are not proven. Rather, they are assumed to be true and used to constrain the solution space for the proof.

Assumptions need to be proven too.

The correctness of a proof depends on the correctness of the assumptions<sup>1</sup> made on the design inputs. Should any assumption be wrong, the proof no longer stands. An assumption on a design's inputs thus becomes an assertion to be proven on the upstream design supplying those inputs.

Semi-formal tools combine property checking with simulation.

Semi-formal tools are hybrid tools that combine formal methods with simulation. Semi-formal tools are used to bridge the gap between the capacity of current formal analysis engines and the size and complexity of the design to be verified. Rather than try to prove all assertions from the reset state, they use intermediate simulation information—such as the current state of a design—as a starting point for proving or disproving assertions.

Use formal methods to prove cases uncovered in simulation.

Formal verification does not replace simulation or make it obsolete. Simulation (including simulated assertions) is the lawnmower of the verification garden: It is still the best technology for covering broad swaths of functionality and for weeding out the easy-to-find

---

1. The formal verification community calls these input assertions “constraints.” I used the term “assumptions” to differentiate them from random-generation constraints, which are randomization concepts.

and some not-so-easy-to-find bugs. Formal verification puts the finishing touch on those hard-to-reach corners in critical and important design sections and ensures that the job is well done. Using functional coverage metrics collected from simulation (for example, request patterns on an arbiter), conditions that remain to be verified are identified. If those conditions would be difficult to create within the simulation environment, it may be easier to prove the correctness of the design for the remaining uncovered cases.

Formal verification should replace ad hoc unit-level verification.

When a designer completes the coding of a design unit—a single or a few modules implementing some elementary function—he or she verifies that it works as intended. This verification is casual and usually the waveform viewer is used to visually inspect the correctness of the response. As mentioned in “Waveform Viewers” on page 39, assertions should be used to specify the signal relationships that define the implementation as “correct” instead of looking for them visually. Once these relationships are specified using assertions, why not try to prove or disprove them using formal technology instead of simulating the design?

Assertion specification is a complex topic.

This simple introduction to assertions does not do justice to the richness and power—and ensuing complexity—of assertions. Entire books have already been written about the subject and should be consulted for more information. Chapter 3 and 7 of the *Verification Methodology Manual for SystemVerilog* provide a lot of guidelines for using assertions with simulation and formal technologies.

---

## REVISION CONTROL

---

Are we all looking at the same thing?

One of the major difficulties in verification is to ensure that what is being verified is actually what will be implemented. When you compile a SystemVerilog source file, what is the guarantee that the design engineer will use that *exact same file* when synthesizing the design?

When the same person verifies and then synthesizes the design, this problem is reduced to that person using proper file management discipline. However, as I hope to have demonstrated in Chapter 1, having the same person perform both tasks is not a reliable functional verification process. It is more likely that separate individuals perform the verification and synthesis tasks.

Files must be centrally managed.

In very small and closely knit groups, it may be possible to have everyone work from a single directory, or to have the design files distributed across a small number of individual directories. Everyone agrees where each other's files are, then each is left to his or her own device. This situation is very common and very dangerous: How can you tell if the designer has changed a source file and maybe introduced a functional bug since you last verified it?

It must be easy to get at all the files, from a single location.

This methodology is not scalable either. It quickly breaks down once the team grows to more than two or three individuals. And it does not work at all when the team is distributed across different physical or geographical areas. The verification engineer is often the first person to face the non-scalability challenge of this environment. Each designer is content working independently in his or her own directories. Individual designs, when properly partitioned, rarely need to refer to some other design in another designer's working directory. As the verification engineer, your first task is to integrate all the pieces into a functional entity. That's where the difficulties of pulling bits and pieces from heterogeneous working environments scattered across multiple file servers become apparent.

### The Software Engineering Experience

HDL models are software projects!

For over 30 years, software engineering has been dealing with the issues of managing a large number of source files, authored by many different individuals, verified by others and compiled into a final product. Make no mistake: Managing a synthesis-based hardware design project is no different than managing a software project.

Free and commercial tools are available.

To help manage files, software engineers use source control management systems. Some are available, free of charge, either bundled with the UNIX operating systems (RCS, CVS, SCCS), or distributed by the GNU project (RCS, CVS) and available in source form at:

`ftp://prep.ai.mit.edu/pub/gnu`

Commercial systems, some very sophisticated, are also available.

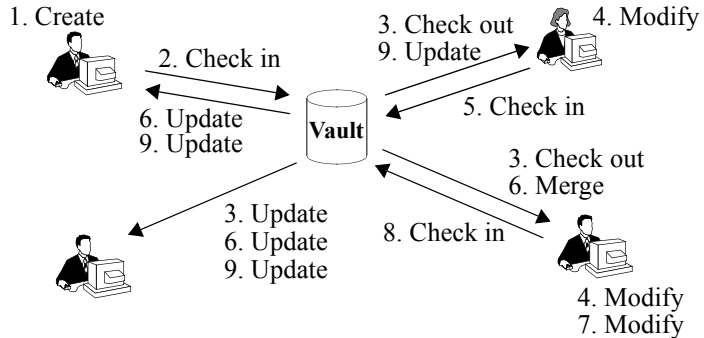
All source files are centrally managed.

Figure 2-17 shows how source files are managed using a source control management system. All accesses and changes to source



files are mediated by the management system. Individual authors and users interact solely through the management system, not by directly accessing files in working directories.

**Figure 2-17.**  
Data flow in a  
source control  
system



The history of a  
file is main-  
tained.

Source code management systems maintain not only the latest version of a file, but also keep a complete history of each file as separate *versions*. Thus, it is possible to recover older versions of files, or to determine what changed from one version to another. It is a good idea to frequently *check in* file versions. You do not have to rely on a backup system if you ever accidentally delete a file. Sometimes, a series of modifications you have been working on for the last couple of hours is making things worse, not better. You can easily roll back the state of a file to a previous version known to work.

The team owns  
all the files.

When using a source management system, files are no longer owned by individuals. Designers may be nominally responsible for various sections of a design, but anyone—with the proper permissions—can make any change to any file. This lets a verification engineer fix bugs found in RTL code without having to rely on the designer, busy trying to get timing closure on another portion of the design. The source management system mediates changes to files either through exclusive locks, or by merging concurrent modifications.


## Configuration Management

Each user works  
from a *view* of  
the file system.

Each engineer working on a project managed with a source control system has a private *view* of all the source files (or a subset thereof) used in the project. Figure 2-18 shows how two users may have two

different views of the source files in the management system. Views need not be always composed of the latest versions of all the files. In fact, for a verification engineer, that would be a hindrance. Files checked in on a regular basis by their authors may include syntax errors, be simple placeholders for future work, or be totally broken. It would be very frustrating if the model you were trying to verify kept changing faster than you could identify problems with it.

**Figure 2-18.**  
User views of  
managed  
source files



Configurations  
are created by  
tagging a set of  
versions.

All source management systems use the concept of symbolic tags that can be attached to specific versions of files. You may then refer to particular versions of files, or set of files, using the symbolic name, without knowing the exact version number they refer to. In Figure 2-18, the user on the left could be working with the versions that were tagged as “ready to simulate” by the author. The user on the right, the system verification engineer, could be working with the versions that were tagged as “golden” by the block-level verification engineer.

Configuration  
management  
translates to tag  
management.

Managing releases becomes a problem of managing tags, which can be a complex task. Table 2-1 shows a list of tags that could be used in a project to identify the various versions of a file as it progresses through the design process. Some tags, such as the “Version\_M.N” tag, never move once applied to a specific version. Others, such as the “Submit” tag, move to newer versions as the development of the design progresses. Before moving a tag, it may be a good idea to leave a trace of the previous position of a tag. One possible mechanism for doing so is to append the date to the tag name. For example, the old “Submit” version gets tagged with the new tag “Submit\_060302” on March 2<sup>nd</sup>, 2006 and the “Submit” tag is moved to the latest version.

**Table 2-1.**  
Example tags  
for release  
management

Tag Name	Description
Submit	Ready to submit to functional verification. Author has verified syntax correctness and basic level of functionality.
Bronze	Passes a basic set of functional testcases. Release is sufficiently functional for integration.
Silver	Passes all functional testcases.
Gold	Passes all functional testcases and meets coding coverage guidelines (requires additional corner-case testcases).
To_Synthesis	Ready to submit to synthesis. Usually matches “Silver” or “Gold”.
To_Layout	Ready to submit to layout. Usually matches “Gold”.
Version_M.N	Version that was manufactured. Matches corresponding “To_Layout” release. Future versions of the same chip will move tags beyond this point.
ON_YYMMDD	Some meaningful release on the specified date.

## Working with Releases

Views can become out-of-date as new versions of files are checked into the source management system database and tags are moved forward.

Releases are specific configurations.

The author of the RTL for a portion of the design would likely always work with the latest version of the files he or she is actively working on, checking in and updating them frequently (typically at relevant points of code development throughout the day and at the end of each day). Once the source code is syntactically correct and its functionality satisfies the designer (by proving all embedded assertions or using a few ad hoc testbenches), the corresponding version of the files are tagged as ready for verification.

Users must update their view to the appropriate release.

You, as the verification engineer, must be constantly on the lookout for updates to your view. When working on a particularly difficult testbench, you may spend several days without updating your view to the latest version ready to be verified. That way, you maintain a consistent view of the design under test and limit changes to the testbenches, which you make. Once the actual verification and debugging of the design starts, you probably want to refresh your view to the latest “ready-to-verify” release of the design before running a testbench.

Update often.

When using a concurrent development model where multiple engineers are working in parallel on the same files, it is important to check in modifications often, and update your view to merge concurrent modifications even more often. If you wait too long, there is a greater probability of collisions that will require manual resolution. The concept of concurrently modifying files then merging the differences sounds impossibly risky at first. However, experience has shown that different functions or bug fixes rarely involve modification to the same lines of source code. As long as the modifications are separated by two or three lines of unmodified code, merging will proceed without any problems. Trust me, concurrent development is the way to go!

You can be notified of new releases.

An interesting feature of some source management systems is the ability to issue email notification whenever a significant event occurs. For example, such a system could send e-mail to all verification engineers whenever the tag identifying the release that is ready for verification is moved. Optionally, the e-mail could contain a copy of the descriptions of the changes that were made to the source files. Upon receiving such an e-mail, you could make an informed decision about whether to update your view immediately.

---

## ISSUE TRACKING

---

All your bug are belong to us!

The job of any verification engineer is to find bugs. Under normal conditions, you should expect to find functional irregularities. You should be *really* worried if no problems are being found. Their occurrence is normal and do not reflect the abilities of the hardware designers. Even the most experienced software designers write code that includes bugs, even in the simplest and shortest routines. Now that we’ve established that bugs *will* be found, how will you deal with them?

Bugs must be fixed.

Once a problem has been identified, it *must* be resolved. All design teams have informal systems to track issues and ensure their resolutions. However, the quality and scalability of these informal systems leaves a lot to be desired.

## What Is an Issue?

Is it worth worrying about?

Before we discuss the various ways issues can be tracked, we must first consider what is an issue worth tracking. The answer depends highly on the tracking system used. The cost of tracking the issue should not be greater than the cost of the issue itself. However, do you want the tracking system to dictate what kind of issues are tracked? Or, do you want to decide on what constitutes a trackable issue, then implement a suitable tracking system? The latter position is the one that serves the ultimate goal better: Making sure that the design is functionally correct.

An issue is *anything* that can affect the functionality of the design:

1. Bugs found during the execution of a testbench are clearly issues worth tracking.
2. Ambiguities or incompleteness in the specification document should also be tracked issues. However, typographical errors definitely do not fit in this category.
3. Architectural decisions and trade-offs are also issues.
4. Errors found at all stages of the design, in the design itself or in the verification environment should be tracked as well.
5. If someone thinks about a new relevant testcase, it should be filed as an issue.

When in doubt, track it.

It is not possible to come up with an exhaustive list of issues worth tracking. Whenever an issue comes up, the only criterion that determines whether it should be tracked, is its effect on the correctness of the final design. If a bad design can be manufactured when that issue goes unresolved, it *must* be tracked. Of course, all issues are not created equal. Some have a direct impact on the functionality of the design, others have minor secondary effects. Issues should be assigned a priority and be addressed in order of that priority.

You may choose not to fix an issue.

Some issues, often of lower importance, may be consciously left unresolved. The design or project team may decide that a particular problem or shortcoming is an acceptable limitation for this particu-

lar project and can be left to be resolved in the next incarnation of the product. The principal difficulty is to make sure that the decision was a conscious and rational one!

### The Grapevine System

Issues can be verbally reported.

The simplest, and most pervasive issue tracking system is the *grapevine*. After identifying a problem, you walk over to the hardware designer's cubicle (assuming you are not the hardware designer as well!) and discuss the issue. Others may be pulled into the conversation or accidentally drop in as they overhear something interesting being debated. Simple issues are usually resolved on the spot. For bigger issues, everyone may agree that further discussions are warranted, pending the input of other individuals. The priority of issues is implicitly communicated by the insistence and frequency of your reminders to the hardware designer.

It works only under specific conditions.

The grapevine system works well with small, closely knit design groups, working in close proximity. If temporary contractors or part-time engineers are on the team, or members are distributed geographically, this system breaks down as instant verbal communications are not readily available. Once issues are verbally resolved, no one has a clear responsibility for making sure that the solution will be implemented.

You are condemned to repeat past mistakes.

Also, this system does not maintain any history. Once an issue is resolved, there is no way to review the process that led to the decision. The same issue may be revisited many times if the implementation of the solution is significantly delayed. If the proposed resolution turns out to be inappropriate, the team may end up going in circles, repeatedly trying previous solutions. Without history, you are condemned to repeat it. There is no opportunity for the team to learn from its mistakes. Learning is limited to individuals, and to the extent that they keep encountering similar problems.

### The Post-It System

Issues can be tracked on little pieces of paper.

When teams become larger, or when communications are no longer regular and casual, the next issue tracking system that is used is the 3M Post-It™ note system. It is easy to recognize at a glance: Every team member has a number of telltale yellow pieces of paper stuck around the periphery of their computer monitor.

If the paper disappears, so does the issue.	This evolutionary system only addresses the lack of ownership of the grapevine system: Whoever has the yellow piece of paper is responsible for its resolution. This ownership is tenuous at best. Many issues are “resolved” when the sticky note accidentally falls on the floor and is swept away by the janitorial staff.
Issues cannot be prioritized.	With the Post-It system, issues are not prioritized. One bug may be critical to another team member, but the owner of the bug may choose to resolve other issues first simply because they are simpler and because resolving them instead reduces the clutter around his computer screen faster. All notes look alike and none indicate a sense of urgency more than the others.
History will repeat itself.	And again, the Post-It system suffers from the same learning disabilities as the grapevine system. Because of the lack of history, issues are revisited many times, and problems are recreated repeatedly.

### **The Procedural System**

Issues can be tracked at group meetings.	The next step in the normal evolution of issue tracking is the procedural system. In this system, issues are formally reported, usually through free-form documents such as e-mail messages. The outstanding issues are reviewed and resolved during team meetings.
Only the biggest issues are tracked.	Because the entire team is involved and the minutes of meetings are usually kept, this system provides an opportunity for team-wide learning. But the procedural system consumes an inordinate amount of precious meeting time. Because of the time and effort involved in tracking and resolving these issues, it is usually reserved for the most important or controversial ones. The smaller, less important—but much more numerous—issues default back to the grapevine or Post-It note systems.

### **Computerized System**

Issues can be tracked using databases.	A revolution in issue tracking comes from using a computer-based system. In such a system, issues must be seen through to resolution: Outstanding issues are repeatedly reported loud and clear. Issues can be formally assigned to individuals or list of individuals. Their resolution need only involve the required team members. The computer-based system can automatically send daily or weekly status reports to interested parties.
--	--

A history of the decision making process is maintained and archived. By recording various attempted solutions and their effectiveness, solutions are only tried once without going in circles. The resolution process of similar issues can be quickly looked-up by anyone, preventing similar mistakes from being committed repeatedly.

But it should not be easier to track them verbally or on paper.

Even with its clear advantages, computer-based systems are often unsuccessful. The main obstacle is their lack of comparative ease-of-use. Remember: The grapevine and Post-It systems are readily available at all times. Given the schedule pressure engineers work under and the amount of work that needs to be done, if you had the choice to report a relatively simple problem, which process would you use:

1. Walk over to the person who has to solve the problem and verbally report it.
2. Describe the problem on a Post-It note, then give it to that same person (and if that person is not there, stick it in the middle of his or her computer screen).
3. Enter a description of the problem in the issue tracking database and never leave your workstation?

It should not take longer to submit an issue than to fix it.

You would probably use the one that requires the least amount of time and effort. If you want your team to use a computer-based issue tracking system successfully, then select one that causes the smallest disruption in their normal work flow. Choose one that is a simple or transparent extension of their normal behavior and tools they already use.

I was involved in a project where the issue tracking system used a proprietary X-based graphical interface. It took about 15 seconds to bring up the entire interface on your screen. You were then faced with a series of required menu selections to identify the precise division, project, system, sub-system, device and functional aspect of the problem, followed by several other dialog boxes to describe the actual issue. Entering the simplest issue took *at least* three to four minutes. And the system could not be accessed when working from home on dial-up lines. You can guess how successful that system was...



Email-based systems have the greatest acceptance.

The systems that have the most success invariably use an e-mail-based interface, usually coupled with a Web-based interface for administrative tasks and reporting. Everyone on your team uses e-mail. It is probably already the preferred mechanism for discussing issues when members are distributed geographically or work in different time zones. Having a system that simply captures these e-mail messages, categorizes them and keeps track of the status and resolution of individual issues (usually through a minimum set of required fields in the e-mail body or header), is an effective way of implementing a computer-based issue tracking system.

---

## METRICS

Metrics are essential management technologies.

Managers love metrics and measurements. They have little time to personally assess the progress and status of a project. They must rely on numbers that (more or less) reflect the current situation.

Metrics are best observed over time to see trends.

Metrics are most often used in a static fashion: “What are the values today?” “How close are they to the values that indicate that the project is complete?” The odometer reports a static value: How far have you travelled. However, metrics provide the most valuable information when observed over time. Not only do you know where you are, but also you can know how fast you are going, and what direction you are heading. (Is it getting better or worse?)

Historical data should be used to create a baseline.

When compared with historical data, metrics can paint a picture of your learning abilities. Unless you know how well (or how poorly) you did last time, how can you tell if you are becoming better at your job? It is important to create a baseline from historical data to determine your productivity level. In an industry where the manufacturing capability doubles every 18 months, you cannot afford to maintain a constant level of productivity.

Metrics can help assess the verification effort.

There are several metrics that can help assess the status, progress and productivity of functional verification. Two have already been introduced: code and functional coverage.

## Code-Related Metrics

Code coverage may not be relevant.

*Code coverage* measures how thoroughly the verification suite exercises the source code being verified. That metric should climb

steadily toward 100 percent over time. From project to project, it should climb faster, and get closer to 100 percent.

However, code coverage is not a suitable metric for all verification projects. It is an effective metric for the smallest design unit that is individually specified (such as an FPGA, a reusable component or an ASIC). But it is ineffective when verifying designs composed of sub-designs that have been independently verified. The objective of that verification is to confirm that the sub-designs are interfaced and cooperate properly, not to verify their individual features. It is unlikely (and unnecessary) to execute all the statements.

The number of lines of code can measure implementation efficiency.

The total *number of lines of code* that is necessary to implement a verification suite can be an effective measure of the effort required in implementing it. This metric can be used to compare the productivity offered by new verification technologies or methods. If they can reduce the number of lines of code that need to be written, then they should reduce the effort required to implement the verification.


Lines-of-code ratio can measure complexity.

The *ratio of lines of code* between the design being verified and the verification suite may measure the complexity of the design. Historical data on that ratio could help predict the verification effort for a new design by predicting its estimated complexity.

Code change rate should trend toward zero.

If you are using a source control system, you can measure the *source code changes* over time. At the beginning of a project, code changes at a very fast rate as new functionality is added and initial versions are augmented. At the beginning of the verification phase, many changes in the code are required by bug fixes. As the verification progresses, the rate of changes should decrease as there are fewer and fewer bugs to be found and fixed. Figure 2-19 shows a plot of the expected code change rate over the life of a project. From this metric, you are able to determine if the code is becoming stable, or identify the most unstable sections of a design.

**Figure 2-19.** Ideal code change rate metric over time



## Quality-Related Metrics

Quality is subjective, but it can be measured indirectly.

Quality-related metrics are probably more directly related with the functional verification than other productivity metrics. Quality is a subjective value, yet, it is possible to find metrics that correlate with the level of quality in a design. This is much like the number of customer complaints or the number of repeat customers can be used to judge the quality of retail services.

Functional coverage can measure testcase completeness.

*Functional coverage* measures the range and combination of input and output values that were submitted to and observed from the design, and of selected internal values. By assigning a weight to each functional coverage metric, it can be reduced to a single functional coverage grade measuring how thoroughly the functionality of the design was exercised. By weighing the more important functional coverage measures more than the less important ones, it gives a good indicator of the progress of the functional verification. This metric should evolve rapidly toward 100 percent at the beginning of the project then significantly slow down as only hard-to-reach functional coverage points remain.

A simple metric is the number of known issues.


The easiest metric to collect is the *number of known outstanding issues*. The number could be weighed to count issues differently according to their severity. When using a computer-based issue tracking system, this metric, as well as trends and rates, can be easily generated. Are issues accumulating (indicating a growing quality problem)? Or, are they decreasing and nearing zero?

Code will be worn out eventually.

If you are dealing with a reusable or long-lived design, it is useful to measure the *number of bugs found during its service life*. These are bugs that were not originally found by the verification suite. If the number of bugs starts to increase dramatically compared to historical findings, it is an indication that the design has outlived its useful life. It has been modified and adapted too many times and needs to be re-designed from scratch. Throughout the normal life

cycle of a reusable design, the number of outstanding issues exhibits a behavior as shown in Figure 2-20.

**Figure 2-20.** Number of outstanding issues throughout the life cycle of a design



Interpreting Metrics

Whatever gets measured gets done.

Because managers rely heavily on metrics to measure performance (and ultimately assign reward and blame), there is a tendency for any organization to align its behavior with the metrics. That is why you must be extremely careful to select metrics that faithfully represent the situation and are correlated with the effect you are trying to measure or improve. If you measure the number of bugs found and fixed, you quickly see an increase in the number of bugs found and fixed. But do you see an increase in the quality of the code being verified? Were bugs simply not previously reported? Are designers more sloppy when writing their code since they'll be rewarded only when and if a bug is found and fixed?

Make sure metrics are correlated with the effect you want to measure.

Figure 2-21 shows a list of file names and current version numbers maintained by two different designers. Which designer is more productive? Do the large version numbers from the designer on the left indicate someone who writes code with many bugs that had to be fixed? Or, are they from a cautious designer who checkpoints changes often?

**Figure 2-21.** Using version numbers as a metric

alu_e.vhd	1.15	cpuif_e.vhd	1.2
alu_rtl.vhd	1.234	cpuif_rtl.vhd	1.4
decoder_e.vhd	1.12	regfile_e.vhd	1.1
decoder_rtl.vhf	1.155	regfile_rtl.vhf	1.7
dpath_e.vhd	1.7	addr_dec_e.vhd	1.3
dpath_rtl.vhd	1.176	addr_dec_rtl.vhd	1.6

On the other hand, Figure 2-22 shows a plot of the code change rate for each designer. What is your assessment of the code quality from

the designer on the left? It seems to me that the designer on the right is not making proper use of the revision control system.

---

**Figure 2-22.**  
Using code  
change rate as  
a metric



## SUMMARY

---

Despite reporting many false errors, linting and other static code checking technologies are still the most efficient mechanism for finding certain classes of problems.

Simulators are only as good as the model they are simulating. Simulators offer many performance enhancing options and the possibility to co-simulate with other languages or simulators.

Assertion-based verification is a powerful addition to any verification methodology. This approach allows the quick identification of problems, where and when they occur.

Verification-specific SystemVerilog features offer an increase in productivity because of their specialization to the verification task and their support for coverage-driven random-based verification.

Use code and functional coverage metrics to provide a quantitative assessment of your progress. Do not focus on reaching 100 percent at all cost. Do not consider the job done when you've reached your initial coverage goals.

Use a source control system and an issue tracking system to manage your code and bug reports.



<http://www.springer.com/978-0-387-29221-2>

Writing Testbenches using SystemVerilog

Bergeron, J.

2006, XXVI, 412 p., Hardcover

ISBN: 978-0-387-29221-2