



# fault: A Python Embedded Domain-Specific Language for Metaprogramming Portable Hardware Verification Components

Lenny Truong<sup>(✉)</sup>, Steven Herbst,  
Rajsekhar Setaluri, Makai Mann, Ross Daly,  
Keyi Zhang, Caleb Donovan, Daniel Stanley,  
Mark Horowitz, Clark Barrett, and Pat Hanrahan

Stanford University, Stanford, CA 94305, USA  
lenny@stanford.edu



**Abstract.** While hardware generators have drastically improved design productivity, they have introduced new challenges for the task of verification. To effectively cover the functionality of a sophisticated generator, verification engineers require tools that provide the flexibility of metaprogramming. However, flexibility alone is not enough; components must also be portable in order to encourage the proliferation of verification libraries as well as enable new methodologies. This paper introduces **fault**, a Python embedded hardware verification language that aims to empower design teams to realize the full potential of generators.

## 1 Introduction

The new golden age of computer architecture relies on advances in the design and implementation of computer-aided design (CAD) tools that enhance productivity [11, 21]. While hardware generators have become much more powerful in recent years, the capabilities of verification tools have not improved at the same pace [12]. This paper introduces **fault**,<sup>1</sup> a domain-specific language (DSL) that aims to enable the construction of flexible and portable verification components, thus helping to realize the full potential of hardware generators.

Using flexible hardware generators [1, 16] drastically improves the productivity of the hardware design process, but simultaneously increases verification cost. A *generator* is a program that consumes a set of parameters and produces a hardware module. The scope of the verification task grows with the capabilities of the generator, since more sophisticated generators can produce hardware with varying interfaces and behavior. To reduce the cost of attaining functional coverage of a generator, verification components must be as flexible as their design

<sup>1</sup> <https://github.com/leonardt/fault>.


counterparts. To achieve flexibility, hardware verification languages must provide the metaprogramming facilities found in hardware construction languages [1].

However, flexibility alone is not enough to match the power of generators; verification tools must also enable the construction of portable components. Generators facilitate the development of hardware libraries and promote the integration of components from external sources. Underlying the utility of these libraries is the ability for components to be reused in a diverse set of environments. The dominance of commercial hardware verification tools with strict licensing requirements presents a challenge in the development of portable verification components. To encourage the proliferation of verification libraries, hardware verification languages must design for portability across verification tools. Design for portability will also promote innovation in tools by simplifying the adoption of new technologies, as well as enable new verification methodologies based on unified interfaces to multiple technologies.

This paper presents **fault**, a domain-specific language (DSL) embedded in Python designed to enable the flexible construction of portable verification components. As an embedded DSL, **fault** users can employ all of Python’s rich metaprogramming capabilities in the description of verification components. Integration with **magma** [15], a hardware construction language embedded in Python, is an essential feature of **fault** that enables full introspection of the hardware circuit under test. By using a staged metaprogramming architecture, **fault** verification components are portable across a wide variety of open-source and commercial verification tools. A key benefit of this architecture is the ability to provide a unified interface to constrained random and formal verification, enabling engineers to reuse the same component in simulation and model checking environments. **fault** is actively used by academic and industrial teams to verify digital, mixed-signal, and analog designs for use in research and production chips. This paper demonstrates **fault**’s capabilities by evaluating the runtime performance of different tools on a variety of applications ranging in complexity from unit tests of a single module to integration tests of a complex design. These experiments leverage **fault**’s portability by reusing the same source input across separate trials for each target tool.

## 2 Design

We had three goals in designing **fault**: enable the construction of flexible test components through metaprogramming, provide portable abstractions that allow test component reuse across multiple target environments, and support direct integration with standard programming language features. The ability to metaprogram test components is a vital requirement for scaling verification efforts to cover the space of functionality utilized by hardware generators. Portability widens the target audience of a reusable component and enhances a design team’s productivity by enabling simple migration to different technologies. Integration with a programming language enables design teams to leverage standard software patterns for reuse as well as feature-rich test automation frameworks.



**Fig. 1.** Architectural overview of the **fault** testing system. In a Python program, the user constructs a **Tester** object with a **magma** **Circuit** and records a sequence of test **Actions**. The compiler uses the action sequence as an intermediate representation (**IR**). Backend targets lower the actions **IR** into a format compatible with the corresponding tool and provide an API to run the test and report the results.

Figure 1 provides an overview of the system architecture. **fault** is a DSL embedded in Python, a prolific dynamic language with rich support for metaprogramming and a large ecosystem of libraries. **fault** is designed to work with **magma** [15], a Python embedded hardware construction language which represents circuits as introspectable Python objects containing ports, connections, and instances of other circuits. While **fault** and **magma** separate the concerns of design and verification into separate DSLs, they are embedded in the same host language for simple interoperability. This multi-language design avoids the complexity of specifying and implementing a single general purpose language without sacrificing the benefits of tightly integrating design and verification code.

To construct **fault** test components, the user first instantiates a **Tester** object with a **magma** circuit as an argument. The user then records a sequence of test actions using an API provided by the **Tester** class. Here is an example of constructing a test for a 16-bit Add circuit:

```
tester = Tester(Add16)
tester.poke(Add16.in0, 3)
tester.poke(Add16.in1, 2)
tester.eval()
tester.expect(Add16.out, 5)
```

The `poke` action (method) sets an input value, the `eval` action triggers evaluation of the circuit (the effects of `poke` actions are not propagated until an `eval` action occurs), and the `expect` action asserts the value of an output. Attributes of the `Add16` object refer to circuit ports by name.

**fault**'s design is based on the concept of staged metaprogramming [20]; the user writes a program that constructs another program to be executed in a subsequent stage. In **fault**, the first stage executes Python code to construct a test specification; the second stage invokes a target runtime that executes this specification. To run the test for the 16-bit Add, the user simply calls a method and provides the desired target:

```
tester.compile_and_run("verilator")
tester.compile_and_run("system-verilog", simulator="iverilog")
```

By applying staged metaprogramming, **fault** allows the user to leverage the full capabilities of the Python host language in the programmatic construction of test components. For example, a test can use a native for loop to construct a sequence of actions using the built-in random number library and integer type:

```
for _ in range(32):
    N = (1 << 16) - 1
    in0, in1 = random.randint(0, N), random.randint(0, N)
    tester.poke(Add16.in0, in0)
    tester.poke(Add16.in1, in1)
    tester.eval()
    tester.expect(Add16.out, (in0 + in1) & N)
```

Python for loops are executed during the first stage of computation and are effectively “unrolled” into a flat sequence of actions. Other control structures such as while loops, if statements, and function calls are handled similarly.

Python’s object introspection capabilities greatly enhance the flexibility of **fault** tests. For example, the core logic of the above test can be generalized to support an arbitrary width Add circuit by inspecting the interface:

```
# compute max value based on port width (length)
N = (1 << len(Add.in0)) - 1
in0, in1 = random.randint(0, N), random.randint(0, N)
tester.poke(Add.in0, in0)
tester.poke(Add.in1, in1)
tester.eval()
tester.expect(Add.out, (in0 + in1) & N)
```

This ability to metaprogram components as a function of the design under test is an essential aspect of **fault**’s design. It allows the construction of generic components that can be reused across designs with varying interfaces and behavior.

**fault**’s embedding in Python’s class system provides an opportunity for reuse through inheritance. For example, a design team could subclass the generic Tester class and add a new method to perform an asynchronous reset sequence:

```
class ResetTester(Tester):
    def __init__(self, circuit, clock, reset_port):
        super().__init__(self, circuit, clock)
        self.reset_port = reset_port

    def reset(self):
        # asynchronous reset, negative edge
        self.poke(self.reset_port, 1)
        self.eval()
        self.poke(self.reset_port, 0)
        self.eval()
        self.poke(self.reset_port, 1)
        self.eval()
```

Combining inheritance with introspection, we can augment the the ResetTester to automatically discover the reset port by inspecting port types:

```

class AutoResetTester(ResetTester):
    def __init__(self, circuit, clock):
        # iterate over interface to find reset (assumes exactly one)
        for port in circuit.interface.ports.values():
            if isinstance(port, AsyncResetN):
                reset_port = port
        super().__init__(self, circuit, clock, reset_port)

```

## 2.1 Frontend: Tester API

**fault**'s Python embedding is implemented by the `Tester` class which provides various interfaces for recording test actions as well as methods for compiling and running tests using a specific target. By using Python's class system to perform a shallow embedding [5], **fault** avoids the complexity of processing abstract syntax trees and simply uses Python's standard execution to construct test components. As a result, programming in **fault** is much like programming with a standard Python library. This design choice reduces the overhead of learning the DSL and simplifies aspects of implementation such as error messages, but comes at the cost of limited capabilities for describing control flow. The **fault** frontend described in this paper focuses on implementation simplicity, but the system is designed to be easily extended with new frontends using alternative embeddings.

**Action Methods.** The `Tester` class provides a low-level interface for recording actions using methods. The basic action methods are `poke` (set a port to a value), `expect` (assert a port equals a value), `step` (invert the value of the clock), `peek` (read the value of a port), and `eval` (evaluate the circuit). The `peek` method returns an object containing a reference to the value of a circuit port in the current simulation state. Using logical and arithmetic operators, the user can construct expressions with this object and pass the result to other actions. For example, to expect that the value of the port `O0` is equal to the inverse of the value of port `O1`, the user would write `tester.expect(circuit.O0, ~tester.peek(circuit.O1))`. The `Tester` provides a `print` action to display simulation runtime information included the peeked values.

**Metaprogramming Control Flow.** Notably absent from the basic method interface described above are control flow abstractions. As noted before, standard Python control structures such as loops and `if` statements are executed in the first stage of computation as part of the metaprogram. However, there are cases where the user intends to preserve the control structure in the generated code, such as long-running loops that should not be unrolled at compile time or loops that are conditioned on dynamic values from the circuit state. For example, consider a `while` loop that executes until it receives a ready signal:

```

# Construct while loop conditioned on circuit.ready.
loop = tester._while(tester.peek(circuit.ready))
loop.expect(circuit.ready, 0) # executes inside loop
loop.step(2)                 # executes inside loop
# Check final state after loop has exited
tester.expect(circuit.count, expected_cycle_count)

```

This logic could not be encoded in the metaprogram, because the metaprogram is evaluated before the test is run, and thus does not know anything about the runtime state of the circuit. To capture this dynamic control flow, the `Tester` provides methods for inserting `if-else` statements, for loops, and while loops. Each of these methods returns a new instance of the current `Tester` object which provides the same API, allowing the user to record actions corresponding to the body of the control construct. The `Tester` class provides convenience functions for using these control structures to generate common patterns, such as `wait_on`, `wait_until_low`, and `wait_until_posedge`.

**Attribute Interface.** While the low-level method interface is useful for writing complex metaprograms, simple components are rather verbose to construct. To simplify the handling of basic actions like `poke` and `peek`, the `Tester` object exposes an interface for referring to circuit ports and internal signals using Python's object attribute syntax. For example, to poke the input port `I` of a circuit with value 1, one would write `tester.circuit.I = 1`. This interface supports referring to internal signals using a hierarchical syntax. For example, referring to port `Q` of an instance `ff` can be done with `tester.circuit.ff.Q`.

**Assume/Guarantee.** The `Tester` object provides methods for specifying assumptions and guarantees that are abstracted over constrained random and formal model checking runtime environments. An *assumption* is a constraint on input values, and a *guarantee* is an assertion on output values. Assumptions and guarantees are specified using Python `lambda` functions that return symbolic expressions referring to the input and output and ports of a circuit. For example, the guarantee `lambda a, b, c: (c >= a) and (c >= b)` states that the output `c` is always greater than or equal to the inputs `a` and `b`. Here is an example of verifying a simple ALU using the assume/guarantee interface:

```
# Configuration sequence for opcode register
tester.circuit.opcode_en = 1
tester.circuit.opcode = 0 # opcode for add (+)
tester.step(2)
tester.circuit.opcode_en = 0
tester.step(2)
# Verify add does not overflow
tester.circuit.a.assume(lambda a: a < BitVector[16](32768))
tester.circuit.b.assume(lambda b: b < BitVector[16](32768))
tester.circuit.c.guarantee(
    lambda a, b, c: (c >= a) and (c >= b)
)
```

Note that this example demonstrates the use of `poke` and `step` to initialize circuits not only for constrained random testing, but also for formal verification.

## 2.2 Actions IR

In using the `Tester` API, users construct a sequence of `Action` objects that are used as an intermediate representation (IR) for the compiler. Basic port action

objects, such as `Poke` and `Expect`, simply store references to ports and values. Control flow action objects, such as `While` and `If`, contain sub-sequences of actions, resulting in a hierarchical data-structure similar to an abstract syntax tree. This view of the compiler internals reveals that the metaphor of *recording actions* is really an abstraction over the construction of program fragments.

### 2.3 Backend Targets

**fault** supports a variety of open-source and commercial backend targets for running tests. A target is responsible for consuming an action sequence, compiling it into a format compatible with the target runtime, and providing an API for invoking the runtime. Targets must also report the result of the test either by reading the exit code of running the process or processing the test output.

**Verilog Simulation Targets.** The **fault** compiler includes support for the open-source Verilog simulators **verilator** [17] and **iverilog** [22], plus three commercial simulators. To compile **fault** programs to a **verilator** test bench, the backend lowers the action sequence into a C++ program that interacts with the software simulation object produced by the **verilator** compiler. For **iverilog** and the commercial simulators, the backend lowers the action sequence into a SystemVerilog test bench that interacts with the test circuit through an `initial` block inside the top-level module. One useful aspect of the SystemVerilog backend is its handling of variations in the feature support of target simulators. For example, the commercial simulators use different commands for enabling waveform tracing and **iverilog** uses a non-standard API for interacting with files. Constrained random inputs are generated using rejection or SMT [9] sampling.

**CoSA.** The CoreIR Symbolic Analyzer (CoSA) is a solver-agnostic SMT-based hardware model checker [13]. **fault**'s CoSA target relies on **magma**'s ability to compile Python circuit descriptions to CoreIR [8], a hardware intermediate representation. CoreIR's formal semantics are based on finite-state machines and the SMT theory of fixed-size bitvectors [3]. **fault** action sequences are lowered into CoSA's custom explicit transition system format (ETS) and combined with the CoreIR representation of the circuit to produce a model. CoSA allows the user to specify assumptions and properties, providing a straightforward lowering of **fault** assumptions and guarantees.

**SPICE.** In addition to being able to test designs with Verilog simulators, **fault** supports analog and mixed-signal simulators. Compared to the traditional approach of maintaining separate implementations for digital and analog tests, this is a significantly easier way to write tests for mixed-signal circuits. Basic actions such as `poke` and `expect` are supported in the SPICE simulation mode, but they are implemented quite differently than they are in Verilog-based tests. Rather than emitting a sequential list of actions in an `initial` block, **fault**

compiles `poke` actions into piecewise-linear (PWL) waveforms. Other actions, such as `expect`, are implemented by post-processing the simulation data.

**Verilog-AMS.** For designs containing a mixture of SPICE and Verilog blocks, `fault` supports testing with a Verilog-AMS simulator. This mode is more similar to running SystemVerilog-based tests than SPICE-based tests. In particular, the test bench is implemented using a top-level SystemVerilog module, meaning that a wide range of actions are supported including loops and conditionals. This is a key benefit of using a Verilog-AMS simulator as opposed to a SPICE simulator.

### 3 Evaluation

To demonstrate `fault`'s capabilities, we evaluate the runtime performance of four different testing tasks from the domain of hardware verification. Each task highlights the utility of `fault`'s portability by reusing the same source input across separate trials of different targets. Due to licensing restrictions, we omit the name of the commercial simulators and replace them with a generic name. The code to reproduce these experiments is available in the artifact.<sup>2</sup> Each experiment involves at least one open-source simulator, but reproducing all the results requires access to commercial simulators.

**CGRA Processing Element Unit Tests.** To demonstrate the capability of `fault` as a tool for writing portable tests for digital verification, Fig. 2 reports the runtime performance of a subset of the `lassen` test suite. `lassen` [19] is an open-source implementation of a CGRA processing element that contains a large suite of unit tests using `fault`. Interestingly, we see comparable performance between `verilator` and `commercial simulator 1`, while `commercial simulator 2` is consistently  $\sim 5x$  slower than the others. One important property of the `lassen` test suite is that it generates a new test bench for each operation and input/output pair. This stresses a simulator's ability to efficiently handle incremental changes, since each invocation involves a new top-level test bench file, but an unchanged design under test.

Test	verilator	commercial sim 1	commercial sim 2
test_unsigned_binary	94.483	88.700	519.079
test_smult	31.439	28.668	170.115
test_fp_binary_op	104.117	91.878	571.759
test_stall	10.424	9.629	56.458

**Fig. 2.** Runtime (s) for unit tests of a CGRA processing element collected with a VM running on an Intel(R) Xeon(R) Silver 4214 CPU @ 2.20 GHz with 256 GB of RAM.

<sup>2</sup> [https://github.com/leonardt/fault\\_artifact/blob/master/README.md](https://github.com/leonardt/fault_artifact/blob/master/README.md).



**SRAM Array.** To demonstrate the capability of **fault** as a tool for writing portable tests for analog and mixed-signal verification, we used **OpenRAM** to generate a 16x16 SRAM and then ran a randomized readback test of the design with SPICE, Verilog-AMS, and SystemVerilog simulators. **OpenRAM** [10] is an open-source memory compiler that produces a SPICE netlist and Verilog model.

The results shown in Fig. 3a reveal two interesting trends. First, as expected, SPICE simulations of the array were significantly slower than Verilog simulations (100-1000x). Since **fault** allows the user to prototype tests with fast Verilog simulations, and then seamlessly switch to SPICE for signoff verification, our tool may reduce the latency in developing mixed-signal tests by orders of magnitude. Second, even for simulations of the same type, there was significant variation in the runtime of different simulators. SPICE simulation time varied by about 2x, while Verilog simulation time varied by about 10x. One of the advantages of using **fault** is that it is easy to switch between simulators to find the one that works best for a particular scenario.

Target	Simulator	Runtime (s)	Lines of Code (LoC)
spice	ngspice	117.660	<b>fault</b> 136
spice	comm sim 1	199.868	Handwritten SPICE 223
spice	comm sim 2	98.043	
system-verilog	iverilog	0.238	Handwritten SystemVerilog 189 and Verilog-AMS
system-verilog	comm sim 1	1.081	
system-verilog	comm sim 2	2.807	
verilog-ams	comm sim 1	228.405	

(a) Runtime using a VM on an Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz with 64GB of RAM.

(b) LoC for **fault** and language-specific implementations of the test.

**Fig. 3.** Results for OpenRAM 16x16 SRAM randomized readback test.

We also looked at the amount of human effort required to use **fault** to implement this test as compared to the traditional approach of writing separate testbenches for each simulation language. Since “human effort” is subjective, we used lines of code as a rough metric, as measured from handwritten implementations of the same test in SystemVerilog, Verilog-AMS, and SPICE. Figure 3b shows the results of this experiment: the **fault**-based approach used 136 LoC as compared to 412 LoC for the traditional approach, a reduction of 3.02x.

**CGRA Integration Test Bench.** To observe how **fault** scales to more complex testing tasks, we report numbers for an integration test of the Stanford Garnet CGRA [18]. This test generates an instance of the CGRA chip, runs a simulation that programs the chip for an image processing application, streams the input image data onto the chip, and streams the output image data to a file. The output is compared to a reference software model. Running the test

took 232 min with the **verilator** target, 185 min with **commercial simulator 1**, and 221 min with **commercial simulator 2**. Leveraging the portability of **fault**-based tests could save up to 47 min in testing time. These results were collected using the same machine as the SRAM experiment (see Fig. 3a).

**Unified Constrained Random and Formal.** To demonstrate the utility of the assume/guarantee interface as a unified abstraction for constrained random and formal verification, we compared the runtime performance of using a constrained random target versus a formal model checker to verify the simple ALU property shown in Sect. 2.1. The first test evaluated the runtime performance of verifying correctness of the property on 100 constrained random inputs versus using a formal model checker. The formal model checker provided a complete proof of correctness using interpolation-based model checking [14] in 1.613 s, while constrained random verified 100 samples in 2.269 s (rejection sampling) and 2.799 s (SMT sampling). The second test injected a bug into the ALU by swapping the opcodes for addition and subtraction. The model checker found a counterexample in 1.154 s with bounded model checking [4], while constrained random failed in 2.947 s (rejection sampling) and 1.230 s (SMT sampling). In both cases the model checker was at least as fast as the constrained random equivalent while providing better coverage in the case of no bug. These results were collected using a MacBook Pro (13-in 2017, 4 Thunderbolt, macOS 10.15.2), with a 3.5 GHz Dual-Core Intel i7 CPU, and 16 GB RAM.

## 4 Related Work

Prior work has leveraged using a generic API to Verilog simulators to build portability into testing infrastructures. The **ChiselTest** library [2] and **cocotb** [7] provide this capability for Scala and Python respectively. Using a generic API offers many of the same advantages with regards to test portability, simplicity, and automation, but the lack of multi-stage execution limits the application to more diverse backend targets such as SPICE simulations and formal model checkers. However, because these libraries interact with the simulator directly, they do allow user code to immediately respond to the simulator state, enabling interactive debugging through the host language. **cocotb** also presents a coroutine abstraction that naturally models the concurrency found in hardware simulation. Future work could investigate using **cocotb** as a runtime target for **fault**'s frontend, enabling a similar concurrent, interactive style of testing. Another interesting avenue of work would be to extend **fault**'s backend targets to support lowering **cocotb**'s coroutine abstraction.

## 5 Conclusion

The ethos of **fault** is to enable the construction of flexible, portable test components that are simple to integrate and scale for testing complex applications.

The ability to metaprogram test components is essential for enabling verification teams to match the productivity of design teams using generators. **fault**'s portability enables teams to easily transition to different tools for different use cases, and enables the proliferation of reusable verification libraries that are applicable in a diverse set of tooling environments.

While **fault** has already demonstrated utility to design teams in academia and industry, there remains a bright future filled with opportunity to improve the system. Extending the assume/guarantee interface to support temporal properties/constraints and leverage compositional reasoning [6] is essential for scaling the approach to more complex systems. Adding concurrent programming abstractions such as coroutines are essential for capturing the common patterns used in the testing of parallel hardware. Using a deep embedding architecture could significantly improve the performance of generating **fault** test benches.

**Funding.** The authors would like to thank the DARPA DSSoC (FA8650-18-2-7861) and POSH (FA8650-18-2-7854) programs, the Stanford AHA and SystemX affiliates, Intel's Agile ISTC, the Hertz Foundation Fellowship, and the Stanford Graduate Fellowship for supporting this work.

## References

1. Bachrach, J., et al.: Chisel: constructing hardware in a scala embedded language. In: 2012 DAC Design Automation Conference, pp. 1212–1221, June 2012. <https://doi.org/10.1145/2228360.2228584>
2. ucb bar: chisel-testers2 (2019). <https://github.com/ucb-bar/chisel-testers2>
3. Barrett, C., Fontaine, P., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB). [www.SMT-LIB.org](http://www.SMT-LIB.org) (2016)
4. Biere, A., Cimatti, A., Clarke, E., Zhu, Y.: Symbolic model checking without BDDs. In: Cleaveland, W.R. (ed.) TACAS 1999. LNCS, vol. 1579, pp. 193–207. Springer, Heidelberg (1999). [https://doi.org/10.1007/3-540-49059-0\\_14](https://doi.org/10.1007/3-540-49059-0_14)
5. Boulton, R.J., Gordon, A., Gordon, M.J.C., Harrison, J., Herbert, J., Tassel, J.V.: Experience with embedding hardware description languages in HOL. In: Proceedings of the IFIP TC10/WG 10.2 International Conference on Theorem Provers in Circuit Design: Theory, Practice and Experience, pp. 129–156. North-Holland Publishing Co., NLD (1992)
6. Clarke, E.M., Long, D.E., McMillan, K.L.: Compositional model checking. In: LICS, pp. 353–362. IEEE Computer Society (1989)
7. cocotb: cocotb (2019). <https://github.com/cocotb/cocotb>
8. Daly, R.: CoreIR: A simple LLVM-style hardware compiler (2017). <https://github.com/rdaly525/coreir>
9. Dutra, R., Bachrach, J., Sen, K.: SMTsampler: efficient stimulus generation from complex SMT constraints. In: 2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 1–8. IEEE (2018)
10. Guthaus, M.R., Stine, J.E., Ataei, S., Chen, B., Wu, B., Sarwar, M.: OpenRAM: an open-source memory compiler. In: 2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 1–6, November 2016. <https://doi.org/10.1145/2966986.2980098>

11. Hennessy, J.L., Patterson, D.A.: A new golden age for computer architecture. *Commun. ACM* **62**(2), 48–60 (2019). <https://doi.org/10.1145/3282307>
12. Lockhart, D., et al.: Experiences building edge TPU with chisel. In: 2018 Chisel Community Conference (CCC) (2018)
13. Mattarei, C., Mann, M., Barrett, C., Daly, R.G., Huff, D., Hanrahan, P.: CoSA: integrated verification for agile hardware design. In: 2018 Formal Methods in Computer Aided Design (FMCAD), pp. 1–5, October 2018. <https://doi.org/10.23919/FMCAD.2018.8603014>
14. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt, W.A., Somenzi, F. (eds.) *CAV 2003*. LNCS, vol. 2725, pp. 1–13. Springer, Heidelberg (2003). [https://doi.org/10.1007/978-3-540-45069-6\\_1](https://doi.org/10.1007/978-3-540-45069-6_1)
15. phanrahan: magma (2019). <https://github.com/phanrahan/magma> (2019)
16. Shacham, O., Azizi, O., Wachs, M., Richardson, S., Horowitz, M.: Rethinking digital design: why design must change. *IEEE Micro* **30**(6), 9–24 (2010). <https://doi.org/10.1109/MM.2010.81>
17. Snyder, W.: Verilator and systemperl. In: North American SystemC Users' Group, Design Automation Conference (2004)
18. StanfordAHA: Garnetflow (2019). <https://github.com/StanfordAHA/GarnetFlow>
19. StanfordAHA: lassen (2019). <https://github.com/StanfordAHA/lassen>
20. Taha, W., Sheard, T.: MetaML and multi-stage programming with explicit annotations. *Theor. Comput. Sci.* **248**(1–2), 211–242 (2000)
21. Truong, L., Hanrahan, P.: A golden age of hardware description languages: applying programming language techniques to improve design productivity. In: Lerner, B.S., Bodík, R., Krishnamurthi, S. (eds.) *3rd Summit on Advances in Programming Languages, SNAPL 2019*, 16–17 May 2019, Providence, RI, USA. *LIPICs*, vol. 136, pp. 7:1–7:21. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2019). <https://doi.org/10.4230/LIPICs.SNAPL.2019.7>
22. Williams, S.: Icarus verilog (2006). <http://iverilog.icarus.com>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

