

Micro-processor verification using a C++11 sequence-based stimulus engine.

DVCON 2017

Stephan Bourduas
(stephan.bourduas@cavium.com)

Chris Mikulis
(chris.mikulis@cavium.com)

March 1, 2017

1. Motivation, Toolchain description.
2. Sequence generator internals (architecture, code examples).
3. Experimental results.
4. Conclusions, Future work.

Introduction

UVM has become the de-facto industry standard for design verification.

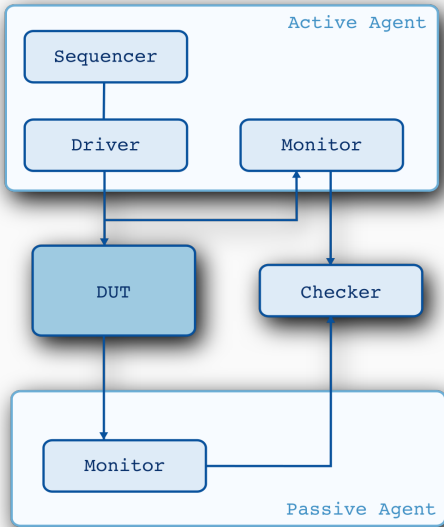
Why didn't we use UVM?

- UVM does not meet **our** stimulus needs for core verification.
 - We decided to create our own sequence-based assembly generator (**SGen**) using new features introduced in C++11.

Why C++11?

- **Lambda** functions enable the creation of closures that capture variables in scope.
- Improved random number generation libraries.
- Polymorphic function wrappers facilitates passing and storing references to **callable** objects.
- Regular expressions.
- The **auto** specifier allows the compiler to deduce types.

Stimulus Generation with UVM



The generic UVM testbench:

- Active agent injects stimulus into the DUT.
- Passive agent monitors output.
- Checker correlates inputs and outputs.

Assumes stimulus is generated by a SV seq/driver pair.

- Not the case for core verification!

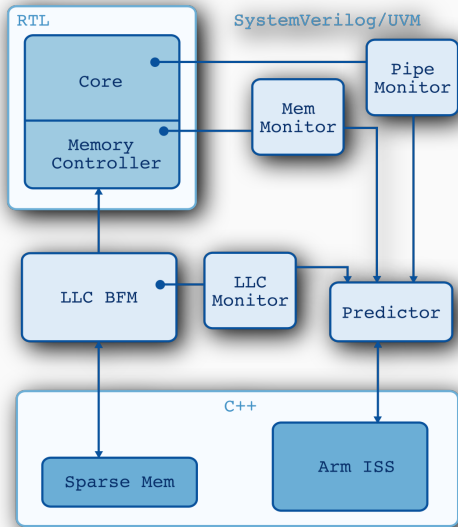
Core Testbench Architecture

Our core testbench:

- Mixed language environment.
- The stimulus is a program binary loaded into memory.
- Core executes program and the BFM reacts.
- UVM seq/drivers used for side-band interfaces.

UVM seq/driver paradigm insufficient:

- We needed to explore alternatives.



Why SGen?

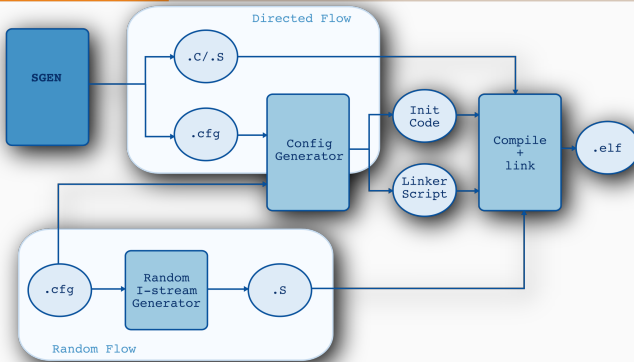
Prior to SGen we had 2 methods to stimulate our design:

1. Hand-written C or assembly tests.
2. Knob-based assembly generator called PPIGen.

The approaches covered opposite ends of the stimulus spectrum:

- Directed tests used for one-off verification tasks and bring up.
- PPIGen is **too** random... controllability decreases as the number of knobs increases.
- **We needed to bridge the gap between directed and fully random stimulus.**

Integrating SGen Into the PPIGen Toolchain



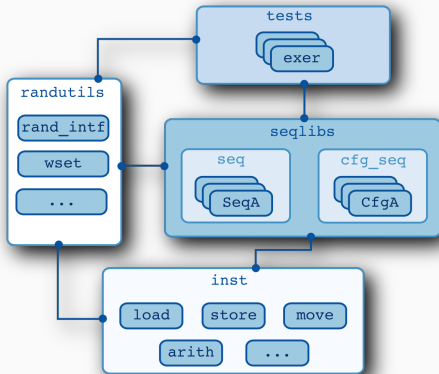
PPIGen

- Directed and random modes.
- Generates init code and linker scripts based on config input (knobs).
- Compiles and link assembly code to generate .elf file.

SGen

- Piggybacks onto the PPIGen directed flow.
- Generates assembly and config files.

Sequence Generator Tool



Library descriptions:

- **inst**: hierarchy of instructions organized by type.
- **seq**: sequences written with specific intent.
 - Includes stimulus and config sequences.
- **tests**: tests achieve goals by using 1 or more sequences.
- **randutils**: provides utility classes that enable randomization.

Randomization Without Constraints

Weighted set:

- Parametrizable container that holds items and their associated weights.
- `pick` method selects a item randomly from the set.
- `pick_and_delete` method selects and removes an item from the set.
- Can contain weighted sets for recursive picking.

Random interface (base class):

- Provides ability to randomize fields belonging to derived classes.
- Maintains a list of `lambda` functions that are executed in fifo order when the `randomize` method is called.
- Lambda functions are added to an object via the `push_callback` method.
- The `push_check` method can be used to install callbacks that check for “constraint violations”.

Randomization — The Weighted Set Class

```

1 // Create 2 weighted sets of chars with 2 items each.
2 // Note that the syntax uses C++11 list initialization .
3 wset<char> w1( {{'a', 100}, {'b', 200}} );
4 wset<char> w2( {{'c', 200}, {'d', 800}} );
5
6 // Create a nested weighted set
7 wset<wset<char>> w;
8
9 // Add items to the weighted set
10 // w2 has a higher weight.
11 w.add_item(w1, 100);
12 w.add_item(w2, 900);
13
14 // Pick a character from nested wset.
15 // Template magic will make the call recursive .
16 char c = w.pick();
  
```

Randomization — The Random Interface Class


```

1 // foo.h
2 class foo : public rand_intf {
3     int x;
4
5     foo() {
6         push_callback(
7             // Set x to 1 by default
8             [this]() { x = 1; }
9         );
10
11        push_check(
12            // make sure x is never
13            // greater than 10;
14            [this]() -> bool { return (x
15                <= 10); }
16        );
17    };
18 };
    
```

```

1 // user_code.cc
2 // Create a weighted set with 2 items
3 // with equal weights.
4 // The set contains an illegal value
5 // of 20 that will cause the
6 // simulation to fail if picked.
7 wset<int> w( {{10, 100}, {20, 100}} );
8
9 auto f = new foo();
10 f->push_callback(
11     // Add another callback to
12     // override the default
13     [&]() { f->x = w.pick(); }
14 );
15
16 // execute all lambda functions in
17 // fifo order.
18 f->randomize();
    
```


Anatomy of a Sequence



- A sequence run generates a **preamble** that sets up base, scratch, index and offset registers and the **main body**.
- The **reg_helper** class provides random register allocation and **reg_init** sequence generates the preamble.
- A **light-weight** sequence is typically short and thus generating a preamble would be wasteful; these sequence rely on the **top-level** sequence for configuration.

Factories

- A hierarchy of factories reflect the class hierarchy in the `inst` package.
- `arm_factory` can instantiate any class that derives from `arm_inst`.
- The `get_types` method returns array of registered type names:
 - Used to instantiate types belonging to a specific factory
 - The user need not know the names of registered types.
 - Facilitates writing generic sequences (new types are automatically used).



Simple Sequence Example

```

1  auto rand_w = wset<unsigned>(1,100); // random number between 1–100
2  auto& f = simd_factory::instance(); // get reference to desired factory
3  auto type_vec = f.get_type_instances(); // return vector of type instances
4
5  // create weighted set of instances with randomized weights.
6  wset<simd_factory::inst_type> inst_wset(type_vec
7      , [&rand_w](){ return rand_w.pick(); });
8
9  // generate random number of instructions.
10 unsigned num = rand_w.pick();
11 for (int n = 0; n < num; ++n) {
12     auto inst = inst_wset.pick();
13     inst->randomize();
14     do_item(*inst);
15 };
    
```

Refer to [listing 3](#) in the paper for a more detailed code example that shows register initialization and instruction randomization.

Experimental Results

Table 1: Incremental compile and link times reported by `unix time` utility.

Modified base class	user time	cpu time
inst	4.5s	1.6s
seq	3.5s	1.3s
test	1.0s	0.5s
all	4.9s	1.8s

- Table 1 shows compile and link times for when files are modified in different packages.
- The `inst` package has the most reverse dependencies.
- The table shows compile and link times for when files are modified in each package.
- The worst case compile time was 5s.

NOTE: Full compile of our testbench took 9m and incremental compile after touching a single SV test took 6m.

Regressing RTL Changes

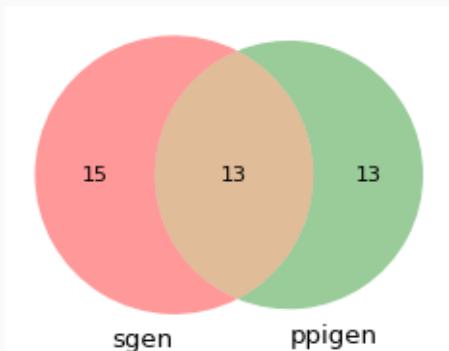
Table 2: Failures broken down by test type for recent regression and exerciser runs.

Series	Tests	Failures		
		Directed	PPIGen	SGen
reg0	421	2	0	6
reg1	446	1	0	15
reg2	417	2	0	29
reg3	410	21	16	34
reg4	388	1	1	8
exer0	47	–	0	5
exer1	46	–	0	6
exer2	41	–	9	26
exer3	35	–	1	0
exer4	55	–	1	0

- RTL designer has added SGen exers to their regression flow.
 - **reg** runs consist of directed tests as well as PPIGen and SGen exers.
 - **exer** runs consist of only random exercisers.
- Directed failures do not necessarily gate check-in.
- SGen has consistently caught bugs that would have escaped standard regression.

PPIGen and SGen Exerciser Runs

Figure 7: Overlap of failure buckets for SGen and PPIGen



- 25k exerciser runs per generator were run over a 1 month period.
- 1731 total failures:
 - SGen accounted for 1560
 - PPIGen accounted for 171
- Once we bin the failures we see that both tools are doing a good job:
 - 41 total buckets with only 13 overlapping.
- SGen is more efficient at hitting bugs despite PPIGen being the more mature tool.

Primary motivation for C++11 was speed:

- Typical exerciser run generates tests with 25k–50k instructions.
- For 500 runs:
 - the average execution time of SGen was 709ms.
 - The number of instructions generated per second was 31k.
- The tool adds no overhead (computes and licenses) to our simulation times.

Conclusion and Future Work

Conclusion

- We bridged the gap between directed and fully random stimulus by creating a sequence-based generator using C++11.
- We were able to express complex dependencies between random variables despite the lack of a constraint engine.
 - Weighted set, random interface and lambda functions.
- New features introduced in C++11 were key enablers.
 - Fast compile and run times increased productivity.
- SGen is currently being used in production to verify the Cavium ThunderX2 core
 - Results show that it is better at uncovering certain types of errors than existing tools.
 - Generated more failures than PPIGen for the same number of runs.

- Continue adding support for SIMD and FP instructions.
- Continue adding to sequence library.
- Improve configuration randomization.
- Explore possibility of using an interpreted language for test writing (i.e. Python front end).

End.