From Gates to Compilers: Putting it All Together

CS250 Laboratory 4 (Version 030316)
Written by Colin Schmidt
Modified by Christopher Yarp
Adapted from Ben Keller

Overview

Up until this point, you have developed and tested the SHA3 accelerator as a separate unit. In this lab, you will connect your accelerator to a fully instantiated RISC-V Rocket processor and write assembly instructions to drive your processor with code instead of a custom testharness. This process will give you a good introduction to the implementation options available to you as you begin your final projects, as well as the process of integrating with the Rocket core.

Deliverables

This lab is more focused on showing you how useful ideas for your projects than having you write your own solutions. As a result the deliverables should follow directly from following the instructions in this document. This lab is due **Thursday, March 10 at 11:59PM**. The deliverables for this lab are:

- (a) written answers to the questions given at the end of this document checked into your git repository as `writeup/lab4-report.pdf` or `writeup/lab4-report.txt`

You are encouraged to discuss your design with others in the class, but you must write your own code and turn in your own work.

Getting lab4

The rocket-chip repository includes several sub-modules, many of which include sub-modules of their own. To avoid cloning the entire RISC-V project, we will avoid simply updating all sub-modules recursively. Follow the procedure below from within your local repository to setup your environment for lab4.

```
git pull template master
git submodule update --init lab4
cd lab4/rocket-chip
git submodule update --init
```

Build Infrastructure

The Lab 4 template files are split into two directories. `lab4/sha3` contains the testharness and build infrastructure to construct your sha3 accelerator as a standalone block. `lab4/rocket-chip` contains the build files for the entire reference chip, with the accelerator source simlinked in. Begin by working in the `lab4` section.

Programming Your Accelerator in RISC-V Assembly

Now that your accelerator is working well on its own, we can begin the process of connecting it to the Rocket core and driving it with actual assembly code. Before we begin to connect any hardware, we want to make sure that we have a good reference for how our accelerator should behave when instructions are passed to it. You will modify `spike`, the RISC-V ISA simulator, so that it understands the custom assembly instruction that we have defined for your accelerator. Then you will write programs that use a combination of C and assembly to drive your accelerator, and run them in `spike` to make sure that the code compiles and behaves as expected.

Rebuilding the ISA Simulator

The current executable for `spike` is configured with the default RISC-V instruction set as defined by the ISA specification. It has also been extended with a custom RoCC instruction called `dummy` that simply reads and writes to memory through the accelerator. You can check that the `spike` executable can handle the dummy instruction by calling it without any additional arguments:

```
% spike --extension=dummy
  usage: spike [host options] <target program> [target options]
  ...
```

It should just print the standard usage information to the command line. However, if you try to load your sha3 acceleration extension, it will complain:

```
% spike --extension=sha3
  unknown extension sha3!
```

You will need to install a new version of `spike` that includes a definition of the `sha3` extension.

From your top-level directory, go to `lab4/riscv-isa-sim`. This directory contains the source code for `spike`. A file defining the functionality of the `sha3` instruction has been implemented for you in `sha3/sha3.h`. Open up this file, and take a look at the implementation.

```
...
class sha3_rocc_t : public rocc_t
{
 public:
  const char* name() { return "sha3"; }

  reg_t custom0(rocc_insn_t insn, reg_t xs1, reg_t xs2)
  {
    switch (insn.funct)
    {
      case 0: // setup: maddr <- xs1; msize <- xs2
...
```

This file should contain C code that defines the sha3 instructions within spikeś infrastructure. To write this type of code you need to explicitly call out accesses to the emulated machine's memory with the functions:

```
p->get_mmu()->load_uint64(reg)
p->get_mmu()->store_uint64(reg,val)
```

Currently this is still a work in progress for sha3 but a fully functional accelerator can be seen in *riscv/dummy-rocc.h*.

This file contains C code that defines the functionality of our new *Custom0* instruction named "sha3". Based on the value of the *funct* field, the code defines two different cases, representing the two sha3 accelerator instructions. In particular, the ADDR instruction demonstrates how to setup accelerator state, and the LEN-START instruction shows how to begin computation.

Although our sha3 accelerator only needs two instructions it would be straightforward to add more instructions, either with more cases for the custom0 instructions or by adding a *Custom1* function. In addition, looking at the dummy accelerator spike implementation may provide other useful examples for RoCC instructions, including ones that return values to the processor through registers rather than through memory.

With the fully defined *Custom0* instruction, you must include the new source file in the existing infrastructure. This involves modifying the `configure` and `configure.ac` files to add sha3 as a default extension when building. These steps have been done for you. You can see the modifications that were made by searching for sha3 in `configure` and `configure.ac` files. Also note that a file called `riscv-sha3.pc.in` is included in the top level spike directory.

Now you should be able to build a new version of `spike` that includes the `sha3` instruction. From the `riscv-isa-sim` directory, run the following commands (make sure to point at your own install path):

```
% export INSTALL_DIR=/scratch/cs250-xx/lab4/toolkit
% mkdir build
```

```
% cd build
% ../configure --prefix=$INSTALL_DIR --with-fesvr=$RISCV
% make
% make install
```

You've installed your new version of `spike` to a directory created in `lab4-ref-chip`. To launch that version of spike, you can either call it explicitly, or add that location to your path:

```
% export PATH=/scratch/cs250-xx/lab4/toolkit/bin:$PATH
```

You can use `which spike` to confirm that you are running the correct version. Running the command `spike --extension=sha3` should no longer throw an error. Note that you must add the above line to your `.bash_profile` or `.bashrc` if you want to automatically point to your new version of `spike` when you start a new terminal session.

Writing C/Assembly Drivers

Now it's time to write some code that exercises the new custom instruction. We have provided several template files that you can use to write programs for your accelerator. They can be found in `lab4/sha3/test`.

Before using the `sha3` instructions, you will need to define a software reference that can check if it is operating correctly. This reference is defined in the file `sha3.h`, notice the function that can check the sha3 hashs.

Now take a look at `sha3-sw.c`. This is the software version of the sha3 acceleration code; it calls the function you just wrote to perform hashes and then "checks" the result with assert statements. Obviously, it is a bit silly to use the same function to calculate and check each value, but this structure parallels the accelerator code that you will write shortly and is useful for direct runtime comparison between the two.

Compile and run the program:

```
% riscv64-unknown-elf-gcc sha3-sw.c -I. -o sha3-sw.rv
% spike pk sha3-sw.rv
```

You'll noticed that we've added an additional `pk` argument to the call to `spike`. `pk` is the RISC-V *proxy kernel*, a sort of lightweight operating system that provides some important features for our C runtime environment. In particular, it provides support for system calls like `assert` and `printf`, so that you can call these functions without having to include extra libraries in your program compilation.

Once you have looked into the software implementation of sha3, open the final C file, `sha3-rocc.c`. This is where you will actually call the `sha3` instruction to evaluate the hashes. This file is structured similarly to the software implementation; it checks the same tests as the previous file. In this case, however, the C code should call the `sha3` instruction instead of the functions defined in `sha3.h`.

Inline assembly instructions in C are invoked with the `asm volatile` command. Before the first instruction, and after each sha3 instruction, the `fence` command is invoked. This ensures that all previous memory accesses will complete before executing subsequent instructions, and is required

to avoid mishaps as the Rocket core and coprocessor pass data back and forth through the shared data cache. (The processor uses the "busy" bit from your accelerator to know when to clear the fence.) A fence command is not strictly required after each custom instruction, but it must stand between any use of shared data by the two subsystems.

The custom commands are invoked with a particular syntax:

```
asm volatile ("custom0 0, %[maddr], %[msize], 0" : :
                          [maddr]"r"(&t1_data), [msize]"r"(t1_size));
```

The `custom0` assembly instruction takes four arguments: *rd*, *rs1*, *rs2*, and *funct*. After the instruction is defined (in quotes), the parameters are passed in after the double colon (`::`). Note that *rs1*, which is an address in memory, should pass in a pointer to an array (rather than passing in a variable directly). Fortunately, you do not need to worry about supplying particular registers to the instruction; the C compiler will make sure that the variables are allocated to registers appropriately.

Read the inline assembly instructions for the test to understand how they work. Once you are finished, compile and run your program:

```
% riscv64-unknown-elf-gcc sha3-rocc.c -I. -o sha3-rocc.rv
% spike --extension=sha3 pk sha3-rocc.rv
```

You are welcome to write a shell script or Makefile to automate the commands to compile and run your code.

Accelerating RISC-V Rocket

Now that you have simulated correct operation of a RISC-V processor with your sha3 accelerator included, you are ready to attach your accelerator as a coprocessor of the Rocket core.

The `rocket-chip` directory contains the code for a chip containing the rocket core along with several other blocks such as caches. Your accelerator has already been symlinked into the `rocket-chip` directory as *sha3*. Since previous labs emulated the RoCC interface, a few modifications to the sha3 project are required. The `test_infrastructure.scala` file defined the RoCC interface. This was required because we did not have the rocket package. Now, the RoCC interface descriptions need to be commented out as they conflict with the description in the rocket package. The rocket package also needs to be imported by `test_infrastructure.scala` and `sha3.scala`. This is accomplished with the line:

```
import rocket._
```

Finally, some of the signal lines for the RoCC interface have changed recently. The `imem` lines have been renamed `autl`.

In addition to the modifications in the sha3 source, there are several files which need to be modified so that the accelerator is included. First, the `Makefrag` file is modified to include sha3 as a add-on module. This line was added into the `Makefrag` after the `CXXFLAGS` line:

```
export ROCKETCHIP_ADDONS ?= sha3
```

Now, the accelerator needs to be wired in. The directory `src/main/scala/` contains the configuration files for the entire rocket-chip. There are several different configurations specified. You will add another set of configurations that attach sha3 to the processor.

Open the file `src/main/scala/PrivateConfig.scala`. This file contains the top-level configuration file for the entire design with sha3. It should look very familiar as it is similar to the way the accelerator configuration is specified which is one of the benefits of jackhammer. The most notable change is a set of lines that tells the chip generator to instantiate a RoCC accelerator:

```
...
case BuildRoCC => Seq(
  RoccParameters(
    opcodes = OpcodeSet.custom0,
    generator = (p: Parameters) => Module(new Sha3Accel()(p))
  )
)
...
```

In the past, only one RoCC accelerator could be connected to rocket-chip at a time. Recent updates to rocket-chip now allow multiple RoCC accelerators to be connected. In your case, we are only instantiating one RoCC accelerator. However, we could instantiate more by simply adding an addition `RoccParameters` object to the sequence following the same pattern used to instantiate sha3. An example of this type of configuration can be found in `Configs.scala`.

The `case RoccMaxTaggedMemXacts => 32` specifies the number of in flight memory requests the rocket-chip will support.

Finally, following lines merge the sha3 configurations with base configurations for generating a CPP emulator, generating a design for ASICs, and generating a design for FPGAs.

```
class Sha3VLSIConfig extends Config(new Sha3Config ++ new DefaultVLSIConfig)
class Sha3FPGAConfig extends Config(new Sha3Config ++ new DefaultFPGAConfig)
class Sha3CPPConfig extends Config(new Sha3Config ++ new DefaultCPPConfig)
```

Now the Rocket processor will connect your accelerator.

RoCC and the C Emulator

To make sure everything has been wired up correctly, we will build the C emulator for the processor. Before we can do this, there are a couple of steps we need to take.

Go into the `emulator` folder and run the following command:

```
% make CONFIG=Sha3CPPConfig
```

Now you should be able to run the same programs that you simulated on the emulated processor:

```
% ./emulator-Top-Sha3CPPConfig pk ../sha3/tests/sha3-sw.rv +dramsim
% ./emulator-Top-Sha3CPPConfig pk ../sha3/tests/sha3-rocc.rv +dramsim
```

This will run considerably slower than `spike`, but should still produce the correct output.

Note that the `printf` and `assert` statements in your C code add considerable overhead to the program execution. There are copies of the two C files named `sha3-sw-bm.c` and `sha3-rocc-bm.c`. These versions have removed all of the `printf` and `assert` statements to speed up the runtime. You can execute these compiled binaries on `spike` to make sure that they run, but no outputs will print to the command line. Instead, you can use these files to benchmark the software implementation of sha3 against your hardware-accelerated implementation.

The following commands will run your benchmark tests on the Rocket C emulator, recording some of the state at each cycle to a file:

```
% ./emulator-Top-Sha3CPPConfig pk ../sha3/tests/sha3-sw-bm.rv +dramsim +verbose \
        3>&1 1>&2 2>&3 | riscv-dis > sha3-sw-bm.out
% ./emulator-Top-Sha3CPPConfig pk ../sha3/tests/sha3-rocc-bm.rv +dramsim +verbose \
        3>&1 1>&2 2>&3 | riscv-dis > sha3-rocc-bm.out
```

Compare the cycle counts reported in the last line of each output file. Feel free to automate this process with a shell script, or by modifying the emulator Makefile.

RTL Simulation

Next you should push your accelerated Rocket processor all the way through the flow. The rocket-chip repository is organized a little differently than the previous labs so there is a different procedure to generate the RTL simulator as well as the synthesized and place-and-routed design.

To simulate the Verilog RTL, execute the following commands from the `rocket-chip` directory.

```
% cd vsim
% make CONFIG=Sha3VLSIConfig
```

To run your program, you can use the following command:

```
% ./simv-Top-Sha3VLSIConfig -q +ntb_random_seed_automatic +dramsim +verbose \
   +max-cycles=100000000 pk ../sha3/tests/sha3-rocc.rv \
   3>&1 1>&2 2>&3 | riscv-dis  > sha3.out
```

This will print the same statements to the terminal and verify that your program still runs correctly.

Synthesis & Place-and-Route

There is no top level Makefile for rocket-chip so you will need to execute each stage of the build process separately. From the rocket-chip directory:

```
% cd vlsi
% make CONFIG=Sha3VLSIConfig srams #this generates the sram models
% cd dc-syn
% make CONFIG=Sha3VLSIConfig
% cd ../icc-par
% make CONFIG=Sha3VLSIConfig
```

Since rocket-chip is complex on its own and also contains Sha3, it will take longer to go through each of these steps.

You have the option of building three different simulators in this directory. You can generate an RTL simulator that mirrors the one you generated in `rocket-chip/vsim`, a post-synthesis simulator, or a post-PAR simulator. For post-synthesis and post-PAR simulation, you should stick to your "benchmark" programs only, as they will take quite some time to complete, and the overhead of `printf` and `assert` statements would slow simulation time even further.

To run a RTL simulation, use the following commands in the `vlsi-sim-rtl` directory.

```
% make CONFIG=Sha3VLSIConfig
% ./simv-Top-Sha3VLSIConfig -q +ntb_random_seed_automatic +dramsim +verbose \
    +max-cycles=100000000 pk ../sha3/tests/sha3-rocc.rv \
    3>&1 1>&2 2>&3 | riscv-dis  > sha3.out
```

To run post-synthesis gate-level simulation, use the following commands in the `vlsi-sim-gl-syn` directory.

```
% make CONFIG=Sha3VLSIConfig
% ./simv-Sha3VLSIConfig -q +ntb_random_seed_automatic +dramsim +verbose \
    +max-cycles=100000000 pk ../sha3/tests/sha3-rocc-bm.rv \
    3>&1 1>&2 2>&3 | riscv-dis  > sha3.out
```

To run post-PAR simulation, use the following commands in the `vlsi-sim-gl-par` directory.

```
% make CONFIG=Sha3VLSIConfig
% ./simv-Sha3VLSIConfig -q +ntb_random_seed_automatic +dramsim +verbose \
    +max-cycles=100000000 pk ../sha3/tests/sha3-rocc-bm.rv \
    3>&1 1>&2 2>&3 | riscv-dis  > sha3.out
```

Submission and Writeup

1. Based on the cycle counts of the two benchmarks, did your coprocessor outperform your software implementation of sha3? By how much? Does this match your expectations?

2. Based on the ICC average power numbers, the clock period, and the number of cycles needed to execute each benchmark, calculate total energy usage in each case. How much energy is saved by using an accelerator?

3. Based on the ICC area report, SHA3 accounts for what percent of the overall design area?

Submission

To complete this lab, you should commit the following files to your private Github repository:

- The files included in the template repository.
- Your answers to the questions above, in a file called `writeup.txt` or `writeup.pdf`.

Some general reminders about lab submission:

- Please note in your writeup if you discussed or received help with the lab from others in the course. This will not affect your grade, but is useful in the interest of full disclosure.
- Please note in your writeup (roughly) how many hours you spent on this lab in total.