

# ECE 6745 Complex Digital ASIC Design, Spring 2025

## Lab 1: ASIC Integer Multiplier

School of Electrical and Computer Engineering  
Cornell University

revision: 2025-02-18-22-22

The first lab assignment is a warmup lab where you will use a standard-cell ASIC toolflow to quantitatively analyze the area, energy, and performance of various integer multiplier designs. You will push three given baseline designs through the ASIC toolflow: a fixed-latency iterative design that always takes the same number of cycles, a variable-latency iterative design that exploits properties of the input operands to reduce the execution time, and a simple single-cycle design. The alternative design is a pipelined integer multiplier that is parameterized by the number of pipeline stages. You are required to implement a register-transfer-level (RTL) model of the alternative design, verify the design using an effective testing strategy, push all designs through the ASIC toolflow, and perform an evaluation comparing the various implementations. Your experience in this lab assignment will create a solid foundation for completing the second lab assignment and the five-week design project.

This handout assumes that you have read and understand the course tutorials and the lab assignment logistics document. To get started, login to an `ecelinux` machine, source the setup script, and clone your lab group's remote repository from GitHub:

```
% source setup-ece6745.sh
% mkdir -p ${HOME}/ece6745
% cd ${HOME}/ece6745
% git clone git@github.com:cornell-ece6745/lab1-groupXX.git
```

where XX is your group number. **You should never fork your lab group's remote repository! If you need to work in isolation then use a branch within your lab group's remote repository.** If you have already cloned your lab group's remote repository, then use `git pull` to ensure you have any recent updates before running all of the tests. You can run all of the tests in the lab like this:

```
% cd ${HOME}/ece6745/lab1-groupXX
% git pull --rebase
% mkdir -p sim/build
% cd sim/build
% pytest ../lab1_imul
```

All of the tests should pass except for the tests related to the multipliers you will be implementing in this lab. For this lab you will be working in the `lab1_imul` subproject which includes the following files in the `sim` subdirectory:

- |                                      |   |
|--------------------------------------|---|
| • <code>IntMulFL.py</code>           | – FL multiplier                               |
| • <code>IntMulFixed.v</code>         | – Verilog RTL fixed-latency multiplier        |
| • <code>IntMulFixed.py</code>        | – Python wrapper for fixed-latency multiplier |
| • <code>IntMulVarCalcShamt.v</code>  | – Verilog RTL model to calculate shift amount |
| • <code>IntMulVarCalcShamt.py</code> | – Python wrapper for calculate shift amount   |

- `IntMulVar.v`
  - `IntMulVar.py`
  - `IntMulScycle.v`
  - `IntMulScycle.py`
  - `IntMulNstage.v`
  - `IntMulNstage.py`
  - `imul-sim`
  - `__init__.py`
  - `test/IntMulFL_test.py`
  - `test/IntMulScycle_test.py`
  - `test/IntMulFixed_test.py`
  - `test/IntMulVarCalcShamt_test.py`
  - `test/IntMulVar_test.py`
  - `test/IntMulNstage_test.py`
  - `test/imul_sim_test.py`
  - `test/__init__.py`
- Verilog RTL variable-latency multiplier
  - Python wrapper for variable-latency multiplier
  - Verilog RTL single-cycle multiplier
  - Python wrapper for single-cycle multiplier
  - Verilog RTL pipelined multiplier
  - Python wrapper for pipelined multiplier
  - Multiplier simulator for evaluation
  - Package setup
  - FL multiplier unit tests
  - RTL single-cycle multiplier unit tests
  - RTL fixed-latency multiplier unit tests
  - RTL calculate shift amount unit tests
  - RTL variable-latency multiplier unit tests
  - RTL pipelined multiplier unit tests
  - Test to make sure multiplier simulator works
  - Test package setup


## 1. Introduction

In ECE 4750, you learned how to evaluate baseline and alternative designs according to several key metrics: execution time (in cycles), area, energy, and cycle time. Using RTL modeling, we were able to quantitatively measure the execution time (in cycles) for a given input dataset for each design, but we relied on qualitative estimates to analyze the area, energy, and cycle time of each design. In this lab, we will learn how to use a standard-cell ASIC toolflow to quantitatively measure the area, energy, and cycle time of a design; we can then combine these measurements with our understanding of the cycle-level performance of each design to conduct a much more detailed design-space exploration.

We have provided you with a functional-level model of an integer multiplier. You can find this functional-level (FL) implementation in `IntMulFL.py` and the associated unit tests in `test/IntMulFL_test.py`. This implementation verifies the functionality and uses a higher-level modeling style compared to register-transfer-level (RTL) modeling. The interface for the FL model and indeed for all of the integer multiplier designs is similar in that they accept two 32-bit numbers and produce a single 32-bit multiplication result. All implementations should treat the input operands and the result as two's complement numbers and thus should be able to handle both signed and unsigned multiplication. The FL model is a simple Python function and the RTL model uses port-based interfaces and the `val/rdy` micro-protocol. By leveraging the `val/rdy` micro-protocol, another module would be able to send request messages to the multiplier and never explicitly be concerned with how many cycles the implementation takes to execute a multiply transaction and return the response message.

## 2. Baseline Design

There are three baseline designs for this lab assignment: a fixed-latency iterative integer multiplier, a variable-latency iterative multiplier, and a simple single-cycle multiplier. The RTL models of each design are provided for you. All baseline designs use the same interface shown in Figure 1.



**Figure 1: RTL Integer Multiplier Interface** – RTL model uses port-based interfaces leveraging val/rdy micro-protocol.

```

1  def imul_fixed_algo( a, b ):
2
3      result = Bits( 32, 0 )
4      for i in range(32):
5          if b[0] == 1:
6              result += a
7              a = a << 1
8              b = b >> 1
9
10     return result

```

**Figure 2: Fixed-Latency Iterative Multiplication Algorithm** – Assumes  $a$  and  $b$  are 32-bit Bits objects. Always uses 32 shifts and subtractions to calculate the partial-products over time. This is executable Python code.

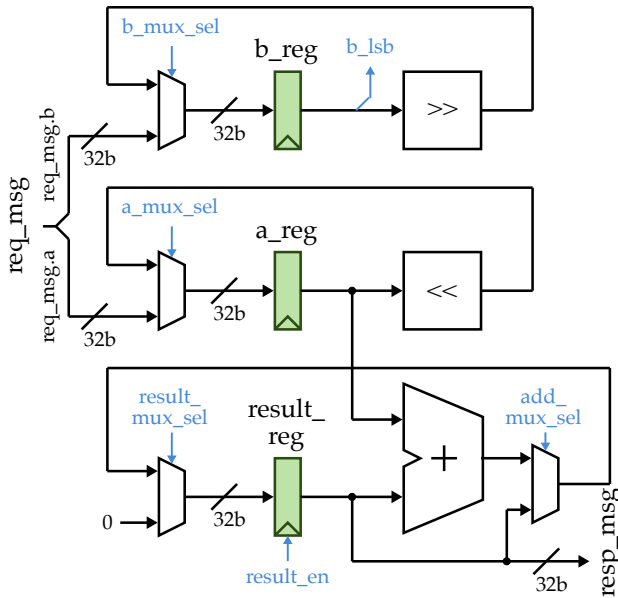
## 2.1. Fixed-Latency Iterative Integer Multiplier

The fixed-latency iterative multiplier will always take approximately 32 cycles. Figure 2 illustrates the fixed-latency iterative multiplication algorithm using “pseudocode” which is really executable Python code. Each iteration we check the least significant bit of the  $b$  operand; if this bit is zero then we shift the  $b$  operand to right and the  $a$  operand to the left, but if this bit is one then we add  $a$  to the result before shifting. Each iteration is essentially calculating a partial product for the multiplication. Note that we say “approximately” because there may be a few extra cycles of overhead in handling the val/rdy micro-protocol (although a more optimized design can remove this overhead).

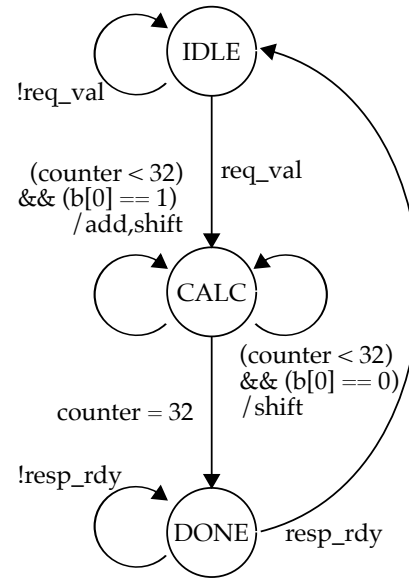
The fixed-latency iterative multiplier RTL model uses a control/datapath split. The datapath for the fixed-latency design is shown in Figure 3. The blue signals represent control/status signals for communicating between the datapath and the control unit. The datapath model uses a structural design style by instantiating a child module for each of the blocks in the datapath diagram. The control unit for the fixed-latency design uses the simple finite-state-machine shown in Figure 4. The IDLE state is responsible for consuming the message from the input interface and placing the input operands into the input registers; the CALC state is responsible for iteratively using adds and shifts to calculate the multiplication; and the DONE state is responsible for sending the message out the output interface. Note that if you implement the FSM exactly as shown each multiply should take 35 cycles: one for the IDLE state, 32 for iterative calculation, one cycle to realize the calculation is done, and one cycle for the DONE state. The extra cycle to realize the calculation is done is because we are not using Mealy transitions from the CALC to DONE states. We need to wait for the counter to reach 32 and then we move into the idle state. The control unit is structured into three parts: a sequential concurrent block for just the state element, a combinational concurrent block for state transitions, and a combinational concurrent block for state outputs. The Verilog code for the fixed-latency iterative multiplier RTL model is in `IntMulFixed.v`.

## 2.2. Variable-Latency Iterative Integer Multiplier

The fixed-latency iterative integer multiplier will always take 35 cycles to compute the result. While it is possible to optimize away the cycle to realize the calculation is done and to eliminate the IDLE/-DONE states, fundamentally this algorithm is limited by the 32 cycles required for the iterative calculation. The variable-latency iterative multiplier takes advantage of the structure in some pairs of input operands to improve performance and energy efficiency. Figure 5 illustrates the variable-latency iterative multiplication algorithm using “pseudocode”. If the  $b$  operand has many consecutive ze-



**Figure 3: Datapath for Fixed-Latency Iterative Integer Multiplier** – All datapath components are 32-bits wide. Shifters are constant one-bit shifters. We use registered inputs with a minimal of logic before the registers.



**Figure 4: Control FSM for Fixed-Latency Iterative Integer Multiplier** – Hybrid Moore/Mealy FSM with Mealy transitions in the CALC state.

ros we don't need to shift one bit per cycle; instead we can shift the B register multiple bits in one step and directly jump to the next required addition. The `calc_shamt` function calculates the shift amount based on the number of trailing zeros in the b operand. Various different implementations of `calc_shamt` are possible: considering more bits will improve the performance but likely increase area and energy.

The variable-latency iterative multiplier RTL model is very similar to the fixed-latency RTL model with a couple key differences. The datapath for the variable-latency design is shown in Figure 6. Notice that we have added a new model that takes the b operand and input and calculates the variable shift amount for both the left and right shifters. The Verilog code for this new model is in `IntMulVarCalcShamt.v`. Refactoring this logic into a separate model enables unit testing the logic before integrating it into the overall design. The control unit for the variable latency design uses the simple finite-state-machine shown in Figure 7. The only difference from the fixed-latency design is that we finish the calculation when the b operand is zero. The Verilog code for the variable-latency iterative multiplier RTL model is in `IntMulVar.v` respectively.

### 2.3. Single-Cycle Integer Multiplier

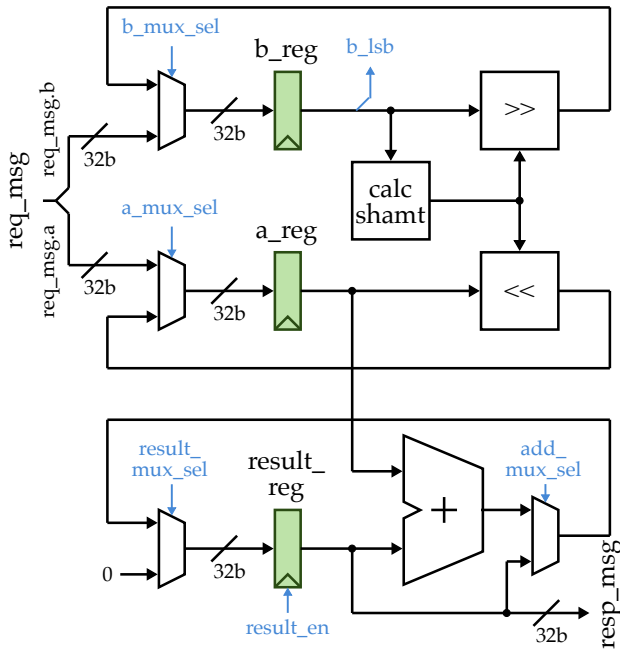
While the iterative multipliers will likely require minimal area, they also require many cycles to calculate each result. As a point of comparison, we will also consider a simple single-cycle integer multiplier. Figure 8 shows the single-cycle RTL model. We use registered inputs for both the operands and the valid bit. If the response interface is not ready, then we stall the multiplier by disabling the register enable signals and combinational propagating the response ready signal to the request ready signal. For the actual multiplier we simply use the `*` operator; we will rely on the ASIC toolflow to choose the most appropriate multiplication hardware. The Verilog code for the single-cycle multiplier RTL model is in `IntMulScycle.v`.

```

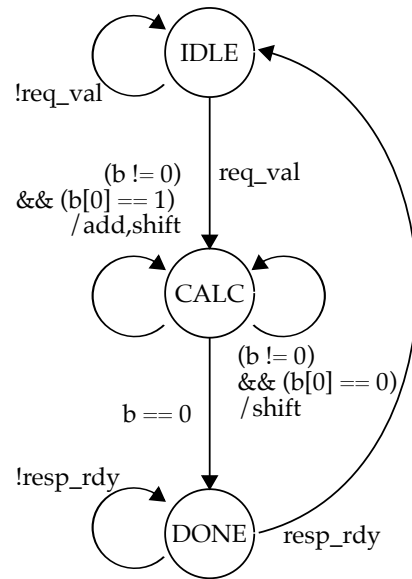
1 def imul_var_algo( a, b ):
2
3     result = Bits( 32, 0 )
4     while b != 0:
5         if b[0] == 1:
6             result += a
7             shamt = calc_shamt( b )
8             a = a << shamt
9             b = b >> shamt
10
11     return result

```


**Figure 5: Variable-Latency Iterative Multiplication Algorithm** – Assumes  $a$  and  $b$  are 32-bit Bits objects. Shifts by more than one to skip over sequences of zeros in the  $b$  operand. Various different  $\text{calc\_shamt}$  functions are possible. This is executable Python code.



**Figure 6: Datapath for Variable-Latency Iterative Integer Multiplier** – All datapath components are 32-bits wide except for the shift amount signal to the variable shifters.



**Figure 7: Control FSM for Variable-Latency Iterative Integer Multiplier** – Hybrid Moore/Mealy FSM with Mealy transitions in the CALC state.



**Figure 8: Single-Cycle Integer Multiplier** – We use registered inputs and combinatorially connect the response ready signal to the request ready and the input register enable signals.

### 3. Alternative Design

For your alternative design, you should implement a pipelined integer multiplier with the same interfaces used in the baseline designs. You should implement an RTL model that is parameterized by the number of pipeline stages; you can assume that the number of stages will be either 1, 2, 4, or 8, although you can also support 16 and 32 stages if you like. As discussed in the tutorials, in this course we will adopt the general policy of using registered inputs for larger blocks, so your pipelined multiplier should start with a set of pipeline registers.

The pipelined integer multiplier RTL model should use a child model that implements a single “partial product step” of the fixed-latency iterative multiplication algorithm. Each partial product step would involve three inputs: one for *a*, *b*, and *result* and three outputs also for *a*, *b*, and *result*. Each partial product step will perform two shifts, an addition, and then mux the result of the addition conditionally based on the least-significant bit of the *b* input. As always, you should unit test the partial product step model, and potentially even build a composition of a few steps for integration testing. Once you are confident in your step model, implement the parameterized RTL model by using elaboration to insert the pipeline registers between the steps. Your pipelined multiplier should always include a total of 32 steps. What changes is just where you put the pipeline registers. If done correctly, your pipelined multiplier should be able to sustain one multiply transaction per cycle.

Note that there are many optimizations you might want to try that may lead to higher quality hardware; you should feel free to explore different approaches if you like. For example, you might want to consider exploring the difference between using structural components within each partial product step vs. doing a flat design without any additional hierarchy within each partial product step. Or you may want to consider exploring the difference between an inelastic vs. elastic pipeline. In an inelastic pipeline, the entire pipeline stalls if the response interface is not ready. In an elastic pipeline, the pipeline can continue flowing even if the response is not ready as long as there are “bubbles” in the pipeline that can be eliminated. Or possibly adding a queue at the response interface can enable better decoupling. Note if students want to explore more complicated multiplier microarchitectures, they should first implement a multiplier using a sequence of partial product steps. Once this first alternative design is working and evaluated, then students can implement another, more complicated alternative design in a different file with different tests.

### 4. Testing Strategy

We have provided you with all of the unit tests for the integer multipliers, although you will need to add your own test scripts for any new models you create. You are also certainly welcome to add more tests for your alternative design if you would like. Most of the unit tests are defined in `test/IntMulFL_test.py`; the other test scripts reuse these unit tests as much as possible. The following commands illustrate how to run all of the tests for the entire project, how to run just the tests for this lab, and how to run just the tests for the FL model, and fixed-latency RTL model. You likely want a test that verifies your pipelined N-stage multiplier is in fact pipelining the multiplications with a cycle-accurate test. Note, you may have to create new testing modules to validate cycle accuracy within a test.

```
% cd ${HOME}/ece6745/lab1-groupXX/sim/build
% pytest ..
% pytest ../lab1_imul
% pytest ../lab1_imul/test/IntMulFL_test.py --verbose
% pytest ../lab1_imul/test/IntMulFixed_test.py --verbose
```

Once your tests are passing for the entire project, you must then generate all of the Verilog test benches which we can then use for four-state RTL simulation, fast-functional gate-level simulation, and back-annotated gate-level simulation.

```
% cd ${HOME}/ece6745/lab1-groupXX/sim/build
% pytest ../lab1_imul/test --test-verilog --dump-vtb
```

## 5. Evaluation

Once you have verified the functionality of your alternative design, you should then use the provided simulator to evaluate the cycle-level performance of the various multiplier designs. You can run the simulator for the first baseline design like this:

```
% cd ${HOME}/ece6745/lab1-groupXX/sim/build
% ../lab1_imul/imul-sim --impl fixed --input small --stats
```

The simulator will display the total number of cycles to execute the specified input dataset. You should study the line traces (with the `--trace` option) and possibly the waveforms (with the `--dump-vcd` option) to understand the reason why each design performs as it does on the various patterns. For your pipelined multiplier, you can use the `--nstages` option to set the number of pipeline stages. Be careful to factor out the pipeline startup overhead for your pipelined multipliers for a fair comparison.

Once you have explored the cycle-level performance of the various multiplier designs, you should push the RTL models through the standard-cell ASIC toolflow to quantify the area, energy, and timing of each design. We would like you to push the following designs through the flow:

- Fixed-latency iterative multiplier
- Variable-latency iterative multiplier
- Single-cycle multiplier
- Pipelined multiplier with one stage
- Pipelined multiplier with two stages
- Pipelined multiplier with four stages
- Pipelined multiplier with eight stages

For each design we would like to analyze the power/energy when executing three specific input datasets: `small` (small random numbers), `large` (large random numbers), and `sparse` (higher probability of zero bits in the input operands). You can generate the Verilog for the fixed latency multiplier using the simulator like this:

```
% cd ${HOME}/ece6745/lab1-groupXX/sim/build
% ../lab1_imul/imul-sim --impl fixed --input small --translate --dump-vtb
% ../lab1_imul/imul-sim --impl fixed --input large --translate --dump-vtb
% ../lab1_imul/imul-sim --impl fixed --input sparse --translate --dump-vtb
```

This will generate the following key files:

- `IntMulFixed__pickled.v`
- `IntMulFixed_imul-fixed-small_tb.v`
- `IntMulFixed_imul-fixed-small_tb.v.cases`
- `IntMulFixed_imul-fixed-large_tb.v`
- `IntMulFixed_imul-fixed-large_tb.v.cases`

- IntMulFixed\_imul-fixed-sparse\_tb.v
- IntMulFixed\_imul-fixed-sparse\_tb.v.cases

A standard-cell ASIC flow can involve many different parameters and a complex sequence of steps. We saw in the tutorials how we can create run scripts to help automate the process, but even writing these scripts can be tedious. So we will be using a simple python script to generate run scripts for us from templates. Each block you wish to push through the standard-cell ASIC flow will require an appropriate YAML file to specify the steps you want to use in the flow along with various parameters. We have provided you seven YAML files that correspond to the seven blocks you will need to push through the flow. These YAML files in the `asic/designs` directory.

- `lab1-fixed.yml`: YAML file for the fixed-latency iterative multiplier
- `lab1-var.yml`: YAML file for the variable-latency iterative multiplier
- `lab1-scycle.yml`: YAML file for the single-cycle multiplier
- `lab1-1stage.yml`: YAML file for the pipelined multiplier with one stage
- `lab1-2stage.yml`: YAML file for the pipelined multiplier with two stages
- `lab1-4stage.yml`: YAML file for the pipelined multiplier with four stages
- `lab1-8stage.yml`: YAML file for the pipelined multiplier with eight stages

You can use `pyhflow` to generate the ASIC flow for the fixed-latency iterative multiplier like this:

```
% cd ${HOME}/ece6745/lab1-groupXX/asic
% mkdir build-lab1-fixed
% cd build-lab1-fixed
% pyhflow ../designs/lab1-fixed.yml
% tree
```

Notice that we use a build directory for pushing a block through the flow. This keeps all of the files associated with each block separate. You can create many build directories, even for the same block as you experiment with different parameters.

You can run the seven steps in sequence like this:

```
% cd ${HOME}/ece6745/lab1-groupXX/asic/build-lab1-fixed
% ./01-synopsys-vcs-rtlsim/run
% ./02-synopsys-dc-synth/run
% ./03-synopsys-vcs-ffglsim/run
% ./04-cadence-innovus-pnr/run
% ./05-synopsys-vcs-baglsim/run
% ./06-synopsys-pt-pwr/run
% ./07-summarize-results/run
```

Although the final summary will display the key metrics you need for your analysis, you should still dig into the various log files and reports at each step. Notice that `pyhflow` also creates a 'run-flow' script which will run all of the steps for you. You can override any parameter in the YAML file from the `pyhflow` command line. So you can rerun the flow for the fixed-latency multiplier with a shorter clock period like this.

```
% cd ${HOME}/ece6745/lab1-groupXX/asic
% pyhflow --clock_period=0.6 ../designs/lab1-fixed.yml
% ./run-flow
```



For the baseline iterative designs you can set the clock constraint to be 0.7 ns, and for the baseline single-cycle design you can set the clock constraint to be 1.5 ns. For your pipelined multiplier you will need to experiment to find a good clock constraint. If you have significant positive slack you are not pushing the tools hard enough. If you have negative slack you are pushing the tools too hard. **Remember that every design must pass all tests using four-state RTL simulation, fast-functional gate-level simulation, and back-annotated gate-level simulation! Remember that every design must meet setup and hold timing after place-and-route!**

Be sure to look at the reports from each run to analyze your results, especially the timing and area reports. Also confirm that your results seem reasonable in case part of the design was aborted by the standard-cell ASIC flow since it can be difficult to see when using the automated scripts. In addition to tables and bar plots, you should include several energy vs. performance plots similar to what we use in lecture to compare the various designs. We recommend having one energy vs. performance plot per input dataset; so plot the results for all of the different microarchitectures executing this dataset on one plot, and then create another plot for a different input dataset. These energy vs. performance plots should enable you to provide deep insight into the various trade-offs in the evaluation section of your report.

## 6. Report

The *Lab Assignment Logistics* document provides general details about the requirements for the report. You must actually read this document to ensure you know what goes in your report and how to format it.

**Alternative Design Section** – In the alternative design section of your report you should briefly discuss how you created a parameterized pipelined multiplier. Be sure to include a diagram illustrating your design. If you also implemented a more complicated multiplier microarchitecture, then you must describe it in the alternative design section.

**Evaluation Section** – In addition to tables and bar plots, you should include several energy vs. performance plots similar to what we use in lecture to compare the various designs. We recommend having one energy vs. performance plot per input dataset; so plot the results for all of the different microarchitectures executing this dataset on one plot, and then create another plot for a different input dataset. These energy vs. performance plots should enable you to provide deep insight into the various trade-offs in the evaluation section of your report. If you also implemented a more complicated multiplier microarchitecture, then you must evaluate it in the alternative design section. **You must include at least one amoeba plot of your alternative design!**

## Acknowledgments

This lab was created by Christopher Batten, Khalid Al-Hawaj, Jason Setter, Derek Lockhart, Shunning Jiang, Yanghui Ou, and Jack Brzozowski as part of the course ECE 6745 Complex Digital ASIC Design at Cornell University.