

LegUp HLS Tutorial for Microsemi PolarFire: Sobel Filtering for Image Edge Detection

Tutorial

December 2020



1 Requirements

Before beginning this tutorial, you should install the following software:

- LegUp HLS 9.0 or later
- Libero SoC 12.5 or later with Modelsim

2 Introduction

This tutorial will introduce you to high-level synthesis (HLS) concepts using LegUp. You will apply HLS to a real problem: synthesizing an image processing application from software written in the C++ programming language. Specifically, you will synthesize a circuit that performs one of the key steps of *edge detection* – a widely used transformation that identifies the edges in an input image and produces an output image showing just those edges. The step you will implement is called *Sobel* filtering. The computations in Sobel filtering are identical to those involved in convolutional layers of a convolutional neural network (CNN). Figure 1 shows an example of Sobel filtering applied to an image. This is the image that will be used as input data in this tutorial.

Sobel filtering involves applying a pair of two 3 x 3 convolutional kernels (also called filters) to an image. The kernels are usually called G_x and G_y and they are shown below in the top of Figure 2. These two kernels “detect” the edges in the image in the horizontal and vertical directions. At each position in the input image, they are applied separately and then combined to produce a pixel value in the output image. The output value is approximated by:

$$G = |G_x| + |G_y|$$

Where the G_x and G_y terms are computed by multiplying each filter value by a pixel value and then summing the products together.



Figure 1: Sobel Filtering Before and After

3x3 convolutional Sobel filters:

-1	0	+1
-2	0	+2
-1	0	+1

Gx

+1	+2	+1
0	0	0
-1	-2	-1

Gy

$$Gx = (-1*3)+(1*4)+(-2*3)+(2*5)+(-1*2)+(1*4) = 7$$

$$Gy = (1*3)+(2*1)+(1*4)+(-1*2)+(-2*3)+(-1*4) = -3$$

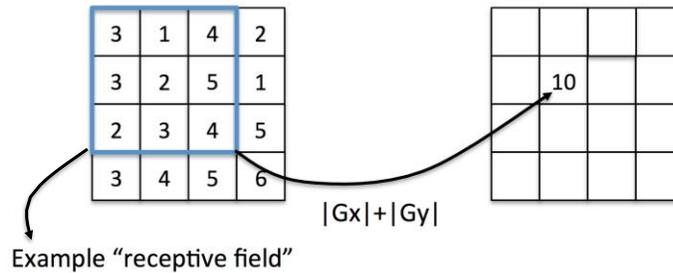


Figure 2: Sobel filters and computational action to compute one pixel in the output image from an input image.

The bottom of Figure 2 shows an input image with 4 rows and 4 columns, where the value in each cell represents a pixel color. The figure illustrates the action of applying the Sobel filters at one position in the input image to compute one value in the output image. The input image pixels involved in the computation are often referred to as the *receptive field*.

A question that may occur to you is: what happens at the edges of the image? i.e., the locations where the placement of the 3 x 3 Sobel filters would "slide off" the edge of the input image. In this tutorial, our program/circuit will simply place 0s into the perimeter of the output image corresponding to such edge locations – this is referred to as *padding* and it is a commonly done technique.

The files for this tutorial can be found on github:

<https://github.com/LegUpComputing/legup-examples/tree/master/tutorials/sobel>

You can download all the files using this link:

<https://github.com/LegUpComputing/legup-examples/archive/master.zip>

Extract this zip file and navigate to the sobel folder:

```
legup-examples-master\tutorials\sobel
```

3 Basic Implementation

In this section, we will use LegUp HLS to compile the Sobel filter to hardware, without any modifications to the C++ code.

First start the LegUp GUI:

On Windows, this can be done by double-clicking on the LegUp shortcut either in the start menu or the desktop.

On Linux, make sure that $\$(LEGUP_INSTALLATION_DIR)/legup/bin$ is on your PATH and the LegUp GUI can be opened by running the following command:

```
legup_hls
```

You will first see a dialog box to select a workspace directory as shown in Figure 3 below. You can use the default workspace for all parts of this tutorial by clicking on *OK*.

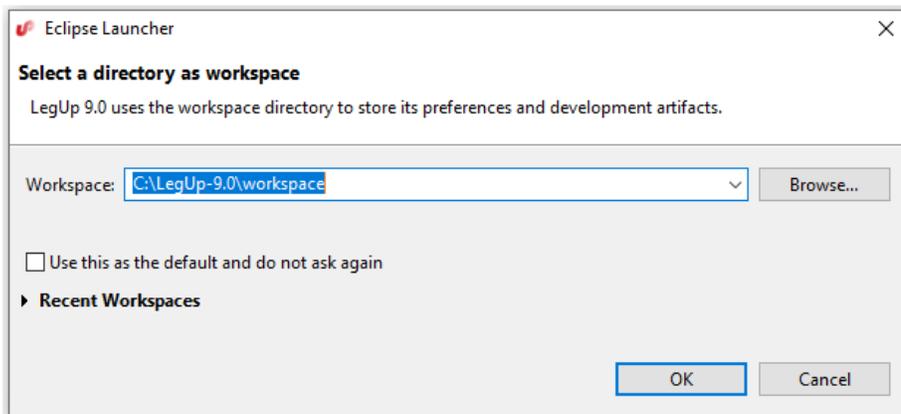


Figure 3: Choosing a LegUp workspace.

Once the LegUp GUI opens, under the File menu, choose New and then LegUp C/C++ project as shown below in Figure 4.

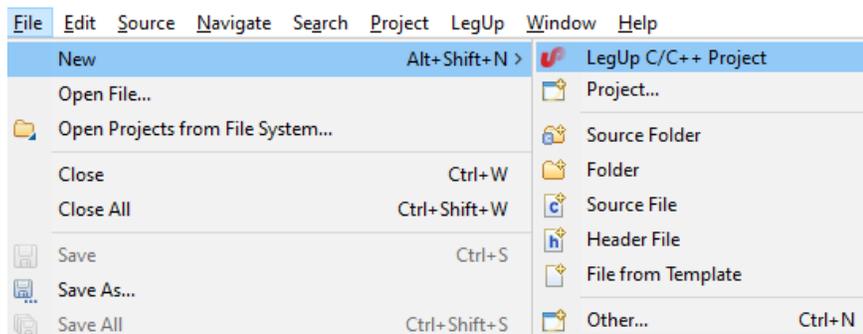


Figure 4: Create a new LegUp C/C++ project

For the project name, enter “sobel_part1” as shown in Figure 5. Then click on *Next*.

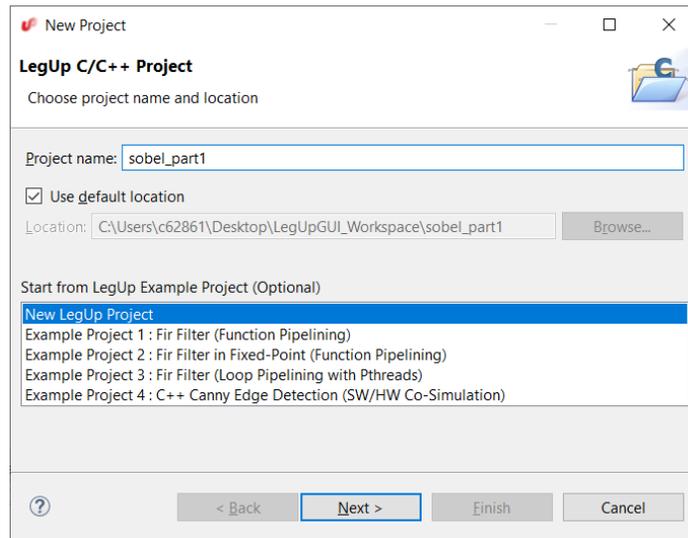


Figure 5: Creating a new LegUp project

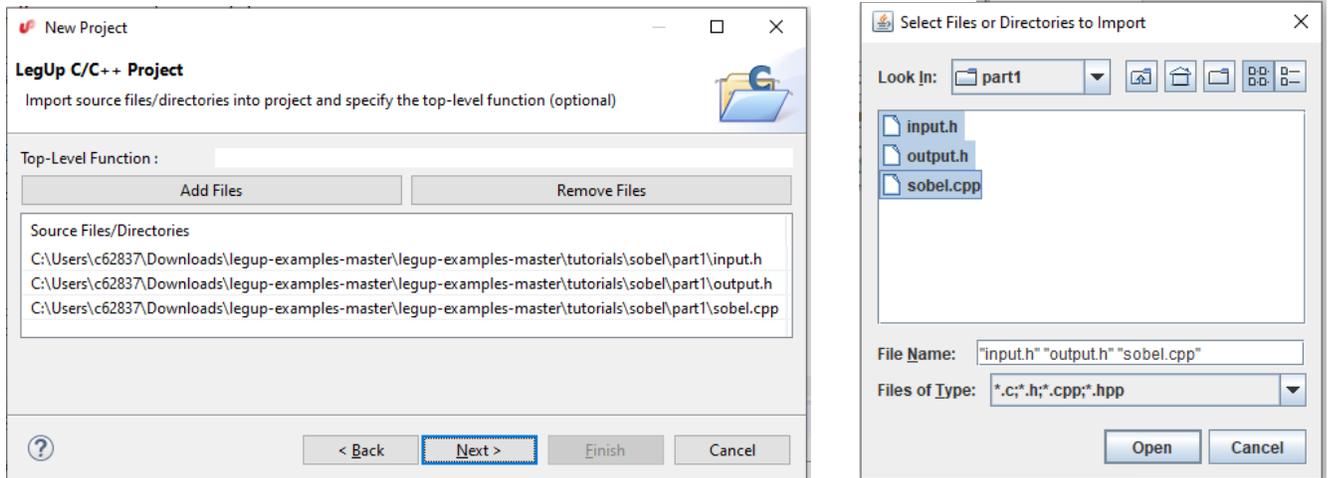


Figure 6: Adding part1 source files into the project.

As shown in Figure 6, you can leave the *Top-Level Function* blank since the `sobel_filter` function in `sobel.cpp` already has a pragma to indicate the top-level function. Now click on *Add Files* to import the source files for part 1 of this tutorial into the project. Navigate to where you have downloaded the tutorial files and go into the *part1* directory. You can hold shift to select all three source files: `input.h`, `output.h` and `sobel.cpp`. After you have added the source files to the project, click on *Next*.

Next you will then see a dialog box where you can specify your own testbench, which are not needed for this part of the tutorial. So, click *Next*, without changing any of the options.

Finally, to complete the project creation, you will choose the FPGA device you intend to target. Use the selections shown in Figure 7, for *FPGA family* choose PolarFire. For *FPGA Device*, you have an option to choose “MPF300TS-1FCG1152I on the MPF300 Board” or use another PolarFire device that is not listed. For this tutorial, we will use another PolarFire device, MPF100T-FCVG484I, which can be used with Microsemi Libero’s free *Silver* license (the bigger MPF300TS-1FCG1152I device requires the paid *Gold* license). To use MPF100T, choose *Custom Device* for the *FPGA Device* field, then type in MPF100T-FCVG484I in the *Custom Device* field. Click on *Finish* when you are done. It may take a few moments to create the project.

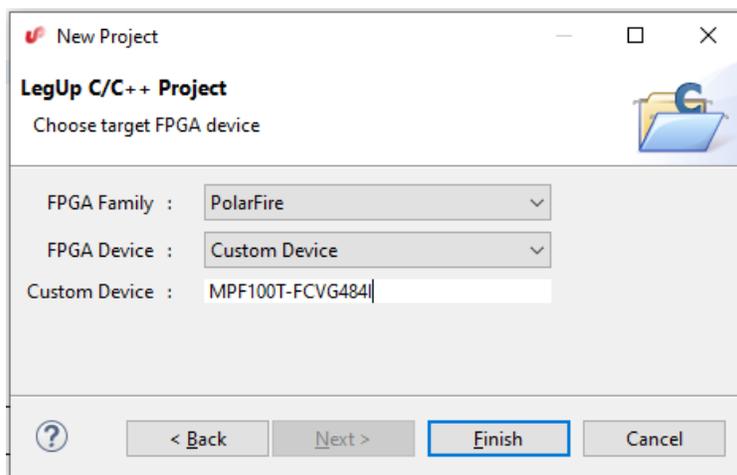


Figure 7: Choose FPGA device

If this is the first time you are using LegUp, you will need to set up the paths to ModelSim (and Microsemi Libero for later parts of this tutorial). To setup the paths, click on *LegUp* on the top menu bar, then click on *Tool Path Settings*. Once the dialog opens, set the paths for *ModelSim Simulator* and *Microsemi Libero SoC* as shown in Figure 8 and click *OK*.

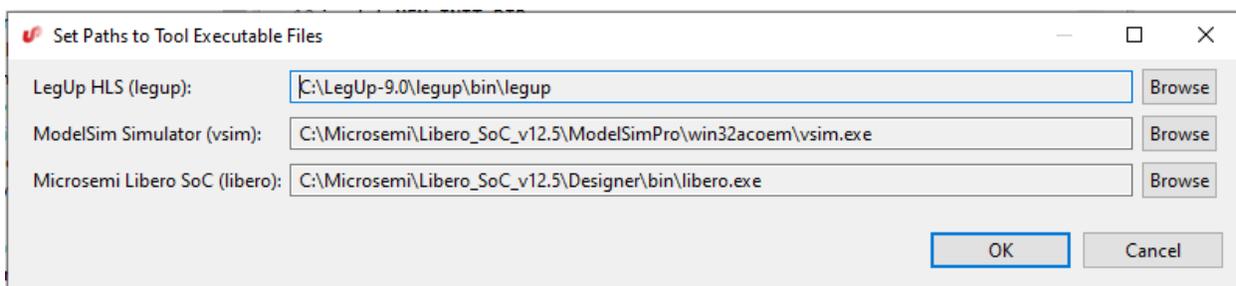


Figure 8: LegUp Tool Path Settings.

An important panel of the GUI is the *Project Explorer* on the left side of the window as shown in Figure 9. This will be used heavily throughout the tutorial to view source files and synthesis reports. Click on the small arrow icon to expand the `sobel_part1` project. You can now double click any of the source files, such as `sobel.cpp`, and you will see the source file appear in the main panel to the right of the *Project Explorer*.

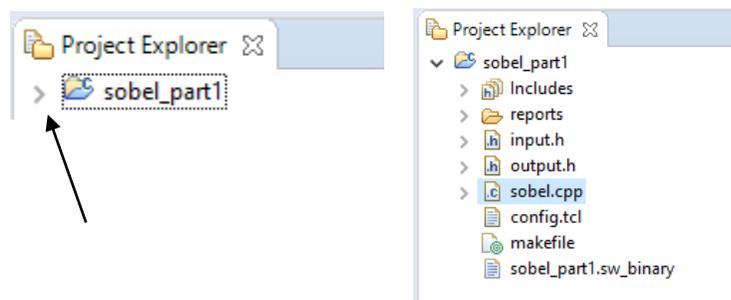


Figure 9: Explorer for browsing source files and reports.

Once the project is created, you should always open one of the source files (such as `sobel.cpp`) or double-click on the `sobel_part1` directory in the Project Explorer pane. This will make `sobel_part1` the active project.

When there are multiple projects open in the workspace, you need to click on the project in the Project Explorer pane or open a file from the project in order to make the project active before running any LegUp HLS commands.

Towards the top of LegUp HLS, you should find a toolbar, as shown in Figure 10, which you can use to execute the main features of the LegUp tool. Hover over each icon to find out their meanings.



Figure 10: LegUp toolbar icons

Starting from the left of Figure 10, the icons are:

- 1) **Add Files to Project**

Then icons for the software development flow:

- 2) **Compile Software with GCC**
- 3) **Run Software that was compiled**
- 4) **Debug Software**
- 5) **Profile Software with gprof**

The hardware development flow icons are:

- 6) **Compile Software to Hardware**
- 7) **Compile Software to Processor/Accelerator SoC**
- 8) **Simulate Hardware**
- 9) **Software/Hardware Co-simulation**
- 10) **Synthesize Hardware to FPGA.**

With the last three icons, you can:

- 11) **Set HLS Constraints**
- 12) **Launch Schedule Viewer**
- 13) **Clean LegUp Project**

These LegUp HLS commands can also be run from the *LegUp* top bar menu.

Figure 11 summarizes the LegUp design flow steps. We initially create the LegUp project and follow a standard software development flow on the C++ (compile/run/debug). Then we apply HLS constraints (i.e., target clock period) and compile the software into hardware using LegUp. We can review reports about the generated hardware. Then we run software/hardware co-simulation to verify the generated hardware. Finally, we can synthesize the hardware to our target FPGA to report the hardware resource usage and Fmax.

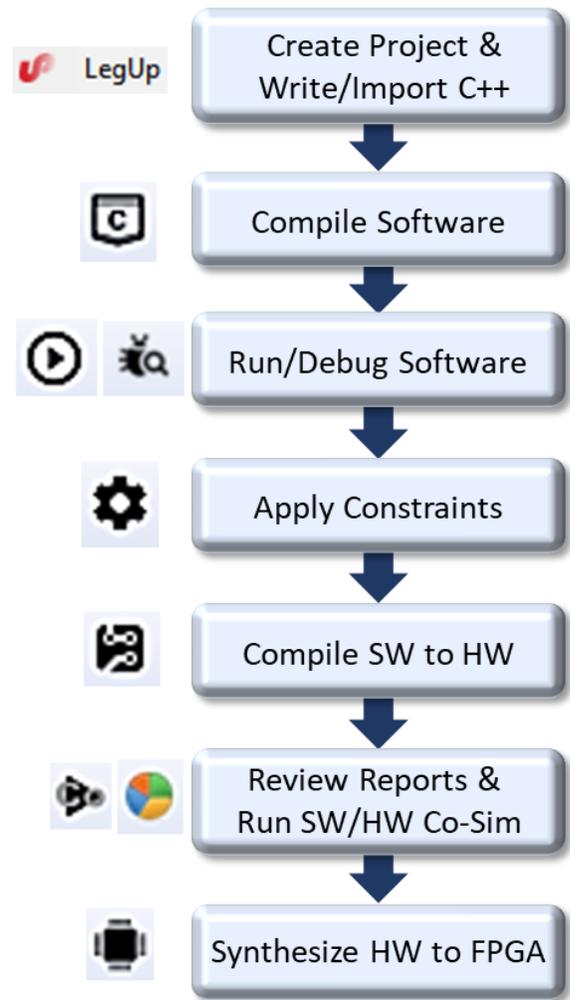


Figure 11: LegUp Design Flow Steps

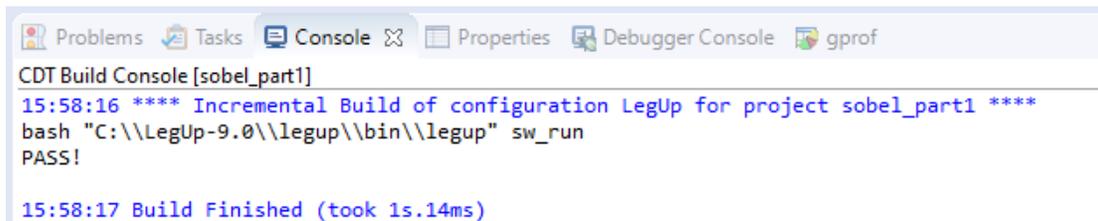
We can now browse through the code in `sobel.cpp`. In the `sobel_filter` function, the first line “`#pragma LEGUP function top`” specifies that the `sobel_filter` function is the top-level function of the project. LegUp will only generate a hardware module for the top-level function and all descendent functions.

The `sobel_filter` top-level function contains a pair of nested loops that iterate over every pixel in the image. Inside the loop body is another pair of nested loops that iterate through the filter window at the current location in the image. For each pixel of the image that is not in the border, the 3x3 area centered on the pixel is convolved with G_x and G_y , then its magnitude summed, to produce the final output pixel.

The `main` function is responsible for verifying the functionality of the `sobel_filter` function. The grayscale (8-bit) input image is stored in the 512x512 array `elaine_512_input` defined in `input.h` and the expected output image is stored in `elaine_512_golden_output` defined in `output.h`. The `main` function passes the input image to the `sobel_filter` function and prints “PASS!” if the computed output image matches the expected output image.

Before compiling to hardware, we should verify that the C++ program is correct by compiling and running the software. This is typical of HLS design, where the designer will verify that the design is functionally correct in software before compiling it to hardware.

Try this by clicking on the *Compile Software* icon in the toolbar (Figure 10). This compiles the software with the GCC compiler. You will see the output from the compilation appearing in the bottom of the screen in the *Console* window of the GUI. Now, execute the compiled software by clicking on the *Run Software* icon in the toolbar. You should see the message *PASS!* appearing in the *Console* window, as shown in Figure 12.



```
Problems Tasks Console Properties Debugger Console gprof
CDT Build Console [sobel_part1]
15:58:16 **** Incremental Build of configuration LegUp for project sobel_part1 ****
bash "C:\\LegUp-9.0\\legup\\bin\\legup" sw_run
PASS!
15:58:17 Build Finished (took 1s.14ms)
```

Figure 12: Console after software execution.

Now we can compile the Sobel filter C++ software into hardware using LegUp by clicking on the toolbar icon to *Compile Software to Hardware*. This command invokes LegUp HLS to compile the top-level `sobel_filter` function into hardware. If the top-level function calls descendant functions, all descendant functions are also compiled to hardware. You can find the generated Verilog code in `sobel_part1.v` as shown in Figure 13.

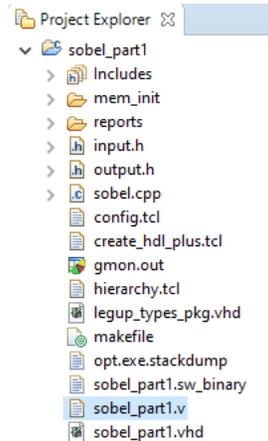


Figure 13: Finding the LegUp-generated Verilog in the Project Explorer.

When the compilation finishes, a LegUp report file (`summary.legup.rpt`) opens. The report shows the RTL interface of the top-level module corresponding to the top-level C++ function, the number of cycles scheduled for each basic block of the function as well as the memories that are used in the hardware. In this example, you will see the top-level RTL module has three interfaces, the standard *Control* interface that is used by any LegUp-generated circuit, and two *Memory* interfaces corresponding to the input and output array arguments of the top-level `sobel_filter` function. In the memory section, there are no memories inside the generated hardware, as the input and output arrays are passed in as arguments into the top-level function. These input/output function arguments are listed as “I/O Memories”.

You can visualize the schedule and control-flow of the hardware using the LegUp schedule viewer GUI. Start the schedule viewer by clicking on the *Launch Schedule Viewer* icon in the toolbar (see Figure 10). In the left panel of the visualizer GUI, you will see the names of the functions and basic blocks of each function. In this example, there is only one function that was compiled to hardware, `sobel_filter`, hence you only see the one function in the Explorer pane on the left.

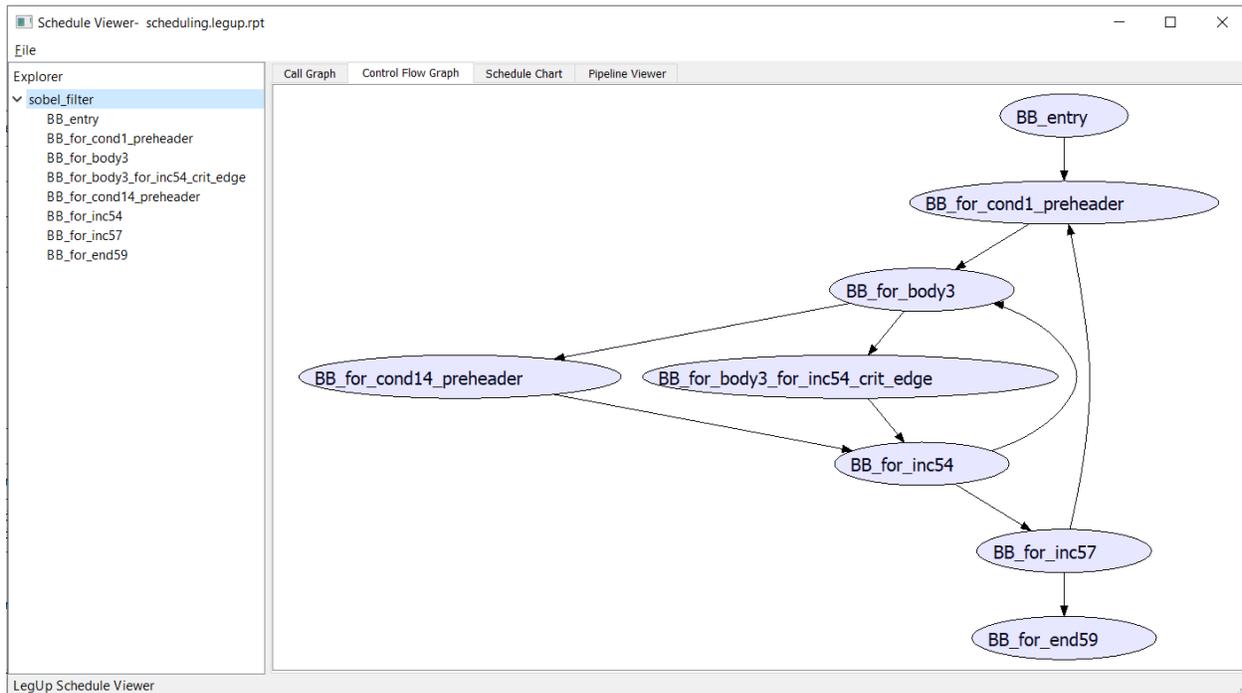


Figure 14: Control-Flow Graph for the Sobel filter.

Double-click on the `sobel_filter` function in the call-graph pane and you will see the control-flow graph for the function, similar to Figure 14. The names of the basic blocks in the program are prefixed with `BB`. Note that the basic block names may be slightly different depending on the version of LegUp you use. The basic block names are not easy to relate back to the original C++ code; however, you can observe that there are two loops in the control-flow graph, which correspond to the two outermost loops in the C++ code for the `sobel_filter` function. The inner loop contains basic blocks: `BB_for_body3`, `BB_for_cond14_preheader`, `BB_for_body3_for_inc54_crit_edge`, and `BB_for_inc54`. Try double clicking on `BB_for_cond14_preheader` (if the basic block names are different from the figure, click on the left-most basic block).

Figure 15 shows the schedule for `BB_for_cond14_preheader`, which is the main part of the inner-most loop body. The middle panel shows the names of the instructions. The right-most panel shows how the instructions are scheduled into states (the figure shows that states 6 to 14 are scheduled for this basic block). Hold your mouse over top of some of the blue boxes in the schedule: you will see the inputs of the current instruction become red and outputs become orange. Look closely at the names of the instructions and try to connect the computations with those in the original C++ program. You will see that there are some loads, additions, subtractions, and shifts.

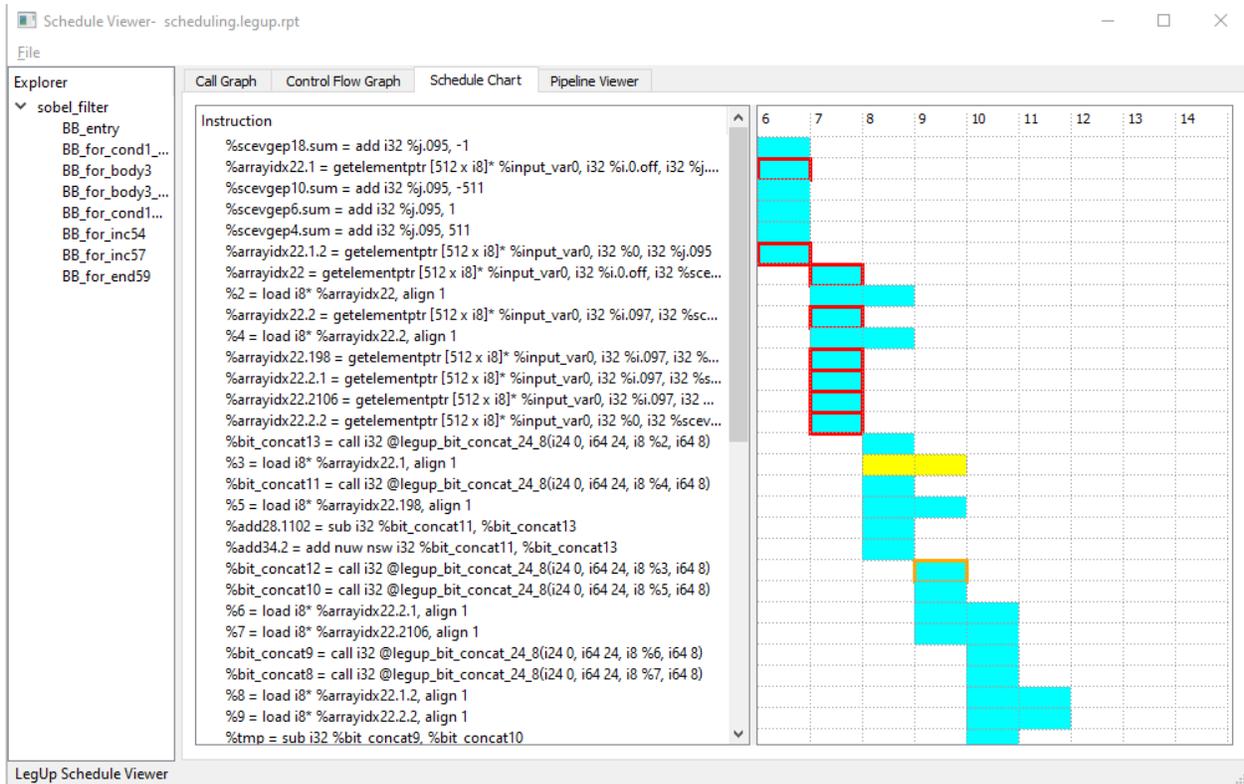


Figure 15: Schedule for the inner-most loop.

Now we can simulate the Verilog RTL hardware with ModelSim to find out the number of cycles needed to execute the circuit – the cycle latency. Close the schedule viewer first, then click the *SW/HW Co-Simulation* icon on the toolbar. SW/HW co-simulation will simulate the generated Verilog module, `sobel_filter_top`, in RTL using ModelSim, while running the rest of the program, `main`, in software. The co-simulation flow allows us to simulate and verify the LegUp-generated hardware without writing a custom RTL testbench.

In the *Console* window, you will see various messages printed by ModelSim related to loading simulation models for the hardware. The hardware may take a few minutes to simulate. We want to focus on the messages near the end of the simulation which will look like this:

```
...
# Cycle latency:      3392549
# ** Note: $finish      : ../simulation/cosim_tb.sv(287)
#   Time: 62730770 ns   Iteration: 1   Instance: /cosim_tb
# End time: 15:27:04 on Dec 08,2020, Elapsed time: 0:00:26
# Errors: 0, Warnings: 0
...
Info: Verifying RTL simulation
...
Retrieving hardware outputs from RTL simulation for sobel_filter
function call 1.
PASS!
...
Number of calls:                1
Cycle latency:      3392549
SW/HW co-simulation: PASS
```

We see that the co-simulation took 3,392,549 clock cycles to finish. The simulation printed “SW/HW co-simulation: PASS!” which indicates that the RTL generated by LegUp matches the software model.

The co-simulation flow uses the return value from the `main` software function to determine whether the co-simulation has passed. If the `main` function returns 0 then the co-simulation will PASS otherwise a non-zero return value will FAIL. Please make sure that your `main` function always follows this convention and returns 0 if the top-level function tests are all successful.

In the `main` function of `sobel_part1`, also called the software testbench, after calling the top-level function, we iterate over every pixel of the computed output image and verify the pixel against the expected value. A mismatch counter is incremented if a pixel is not as expected and this counter value is returned by the `main` function. If all values match, then the `main` function will return 0. Therefore, since co-simulation printed PASS (`main` returned 0) we have verified the generated hardware is correct.

We can also run co-simulation and launch Modelsim to show the Waveforms. From the *LegUp* top menu, select *SW/HW Co-Simulation with Waveforms* as shown in Figure 16.

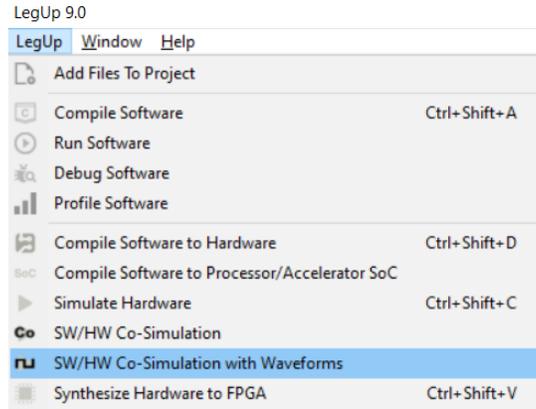


Figure 16: Run SW/HW Co-Simulation with Waveforms

When Modelsim opens it will prompt “Are you sure you want to finish?”. Select “No”. Then you can view the signal waveforms as shown in Figure 17.

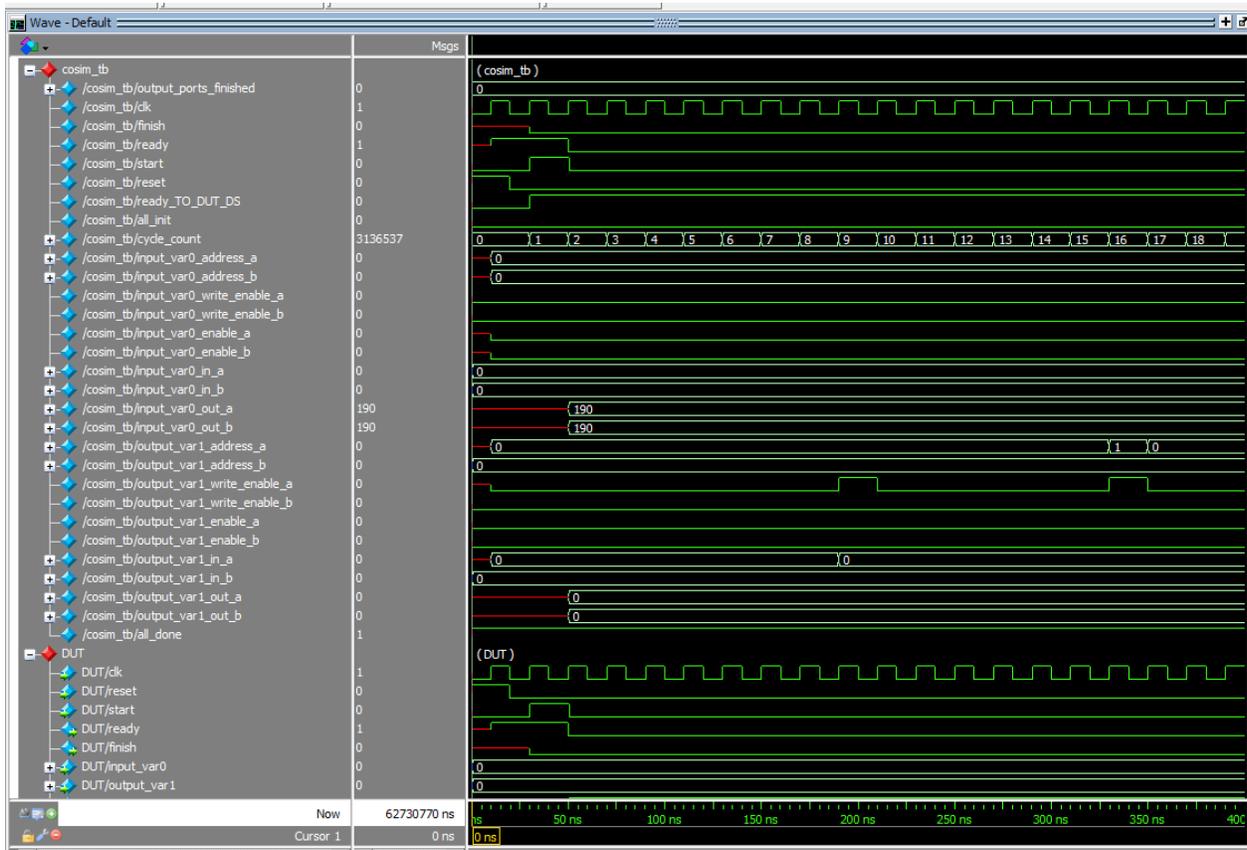


Figure 17: Modelsim waveforms shown during SW/HW Co-Simulation

Libero is the name of Microsemi’s synthesis, placement, routing, and timing analysis tool. LegUp can execute Libero to synthesize, place and route the Verilog to the Microsemi PolarFire FPGA to obtain information such as the resource usage and the Fmax of this design (i.e. the clock period). To do this, click the icon on the toolbar to *Synthesize Hardware to FPGA*. This will automatically invoke Libero to create a Libero project and synthesize the LegUp design targeting the PolarFire FPGA device. Libero may take a while to finish.

Once the command completes, LegUp will open the summary.results.rpt report file. LegUp will summarize the resource usage and Fmax results reported by Libero after place and route. You should get similar results as what is shown below. Your numbers may differ slightly, depending the version of LegUp and Libero you are using. This tutorial used Libero SoC v12.5. The timing results and resource usage might also differ depending on the random seed used in the synthesis tool flow.

===== 2. Timing Result =====

Clock Domain	Target Period	Target Fmax	Worst Slack	Period	Fmax
clk	10.000 ns	100.000 MHz	7.635 ns	2.365 ns	422.833 MHz

===== 3. Resource Usage =====

Resource Type	Used	Total	Percentage
4LUT	1011	108600	0.93
DFF	487	108600	0.45
I/O Register	0	852	0.00
User I/O	0	284	0.00
uSRAM	0	1008	0.00
LSRAM	0	352	0.00
Math	0	336	0.00

Wall-clock time is one of the key performance metrics for an FPGA design, computed as the product of the cycle latency and the clock period. In this case, our cycle latency was 3,392,549 and the clock period was 2.365 ns. The wall-clock time of our implementation is therefore $3,392,549 \times 2.365 \text{ ns} = 8.023 \text{ ns}$.

4 Loop Pipelining

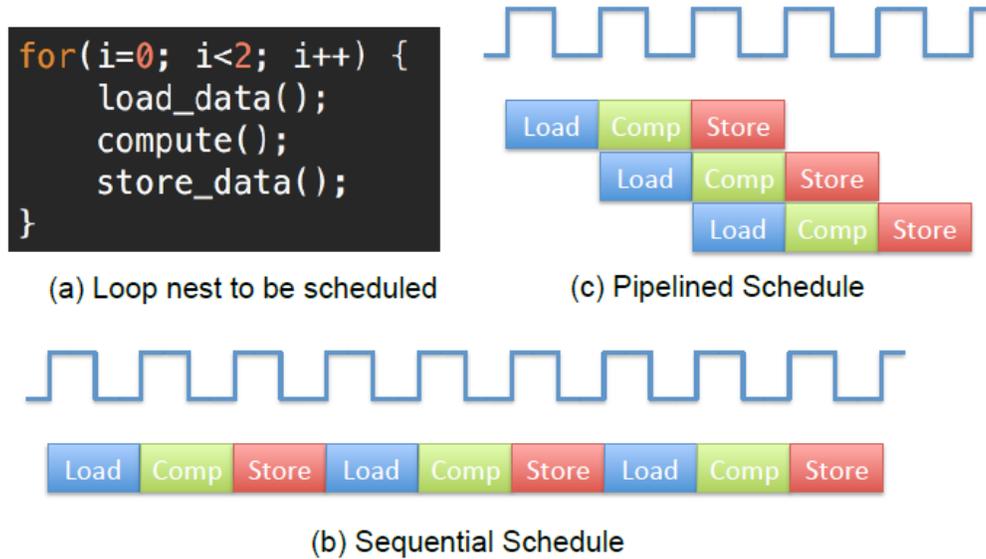


Figure 18: Loop pipelining Example.

In this section, you will use loop pipelining to improve the throughput of the hardware generated by LegUp. Loop pipelining allows a new iteration of the loop to be started before the current iteration has finished. By allowing the execution of the loop iterations to be overlapped, higher throughput can be achieved. The amount of overlap is controlled by the initiation interval (II). The II indicates how many cycles are required before starting the next loop iteration. Thus, an II of 1 means a new loop iteration can be started every clock cycle, which is the best we can achieve. The II needs to be larger than 1 in other cases, such as when there is a resource contention (multiple loop iterations need the same resource in the same clock cycle) or when there are loop-carried dependencies (the output of a previous iteration is needed as an input to the subsequent iteration). Resource contention commonly happens with memory accesses to dual-port block RAMs which can only perform two memory accesses per cycle.

Figure 18 shows an example of loop pipelining. Figure 18(b) shows the sequential loop, where a new loop iteration can start every 3 clock cycles (II=3), and the loop takes 9 clock cycles to finish the final write. Figure 18(c) shows the pipelined loop. In this example, there are no resource contentions or data dependencies. Therefore, the pipelined loop can start a new iteration every clock cycle (II=1) and takes only 5 clock cycles to finish the final write. As shown in this example, loop pipelining can significantly improve the performance of your circuit, especially when there are no data dependencies or resource contentions.

Follow the same procedure you used in the previous part of this tutorial (include the source files for part 2 and target MPF100T-FCVG484I), create a new LegUp project for part 2. Once the project is created, open sobel.cpp from part 2. With the Sobel filter, since each pixel of the output is dependent only on the input image and the constant matrices G_x and G_y , we would like to pipeline the calculation of each pixel.

Open the part 2 source file `sobel.cpp`. The loop pipeline pragma in front of the loop, “`#pragma LEGUP loop pipeline`”, tells LegUp to pipeline the loop and take advantage of loop parallelism:

```
#pragma LEGUP loop pipeline
    for (int i = 0; i < (HEIGHT - 2) * (WIDTH - 2); i++) {
```

You will notice that the pair of nested loops in part 1 (previously using `i` and `j`) have been flattened into one for loop. We have flattened the nested loops to allow us to apply loop pipelining on the entire loop body. Otherwise when we apply loop pipelining on a nested outer loop, LegUp will automatically unroll the inner loops, which would not be possible with 512 iterations.

Now you can synthesize the design by clicking the *Compile Software to Hardware* icon in the toolbar. In the *Console* window, you should see messages like the following:

```
Info: Resource constraint limits initiation interval to 4.
Resource 'input_var0_shared_local_memory_port' has 8 uses per cycle but only 2 units available.
+-----+-----+-----+
| Operation                               | Location                               | Competing Use Count |
+-----+-----+-----+
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 1                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 2                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 3                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 4                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 5                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 6                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 7                    |
| 'load' (8b) operation for array 'input' | line 28 of sobel.cpp                  | 8                    |
+-----+-----+-----+
|                                         | Total # of Competing Uses | 8                    |
+-----+-----+-----+
```

These messages indicate that LegUp cannot achieve an II of 1 (highest throughput) due to resource conflicts – there are 8 loads from the same RAM memory in the loop body. Since RAM blocks are dual ported on an FPGA, we need 4 cycles to perform 8 loads. Therefore, LegUp needs to schedule 4 cycles between successive iterations of the loop (an initiation interval of 4).

A good way to visualize this is to use the LegUp schedule viewer. Click on the *Launch Schedule Viewer* icon again. Double-click on `sobel_filter`, then in the Control Flow Graph, you will see a basic block called `BB_for_body`. Double-click `BB_for_body` to reveal the loop pipeline schedule, similar to that shown in Figure 19. Horizontally from left to right shows the operations performed on successive clock cycles and vertically going down shows successive loop iterations. Here, you can see that the II of the loop is 4 and that a new loop iteration starts every 4 cycles. You can also see the instructions that are scheduled in each cycle for each loop iteration. All instructions that are shown in the same column are executed in the same cycle.

The dark black rectangle on the right illustrates what the pipeline looks like in steady state. In steady state, three iterations of the loop are “in flight” at once. In steady state, you can see that there are two loads in cycle 8 from iteration 1 (the second row down), two loads in cycle 9 from iteration 2 (third row down), and two loads in cycle 11 from iteration 2. Thus the 8 loads are spread out over 4 cycles, making the II = 4.

Cycle:	8	9	10	11
Pipeline Stage:	3	3	3	3
Iteration: 0	%add42 = add ns %bit_select = c %cmp43 = icmp s	%conv48 = selec store i8 %conv4		
Iteration: 1	%bit_concat12 = %bit_concat10 = %4 = load i8* % %5 = load i8* %	%bit_concat9 = %bit_concat8 = %6 = load i8* % %7 = load i8* % %tmp = sub i32 %bit_select4 = %bit_concat5 = %add21.2.1 = su %add27.2102 = s	%bit_concat7 = %bit_concat6 = %add21.299 = ad %add21.2.2 = ad %bit_select2 = %tmp3 = sub i32 %bit_select3 = %bit_concat = c %add27.1.2 = su %add27.2.2 = ad %bit_select1 =	%sub = sub nsw %cond35 = selec %sub38 = sub ns %cond41 = selec
Iteration: 2	%x.086 = phi i3 %y.085 = phi i3 %i.084 = phi i3 tail call void %cmp1 = icmp eq %bit_concat14 = %cond = add nsw	%add14 = add ns %add15 = add ns %add14.2 = add %add15.2 = add %arrayidx50 = g	%arrayidx16 = g %0 = load i8* % %arrayidx16.1 = %arrayidx16.2 = %2 = load i8* % %arrayidx16.187 %arrayidx16.2.1	%bit_concat13 = %1 = load i8* % %bit_concat11 = %3 = load i8* % %add21.191 = su %add27.2 = add

Figure 19: Loop Pipeline Schedule for the Sobel filter.

Now, exit the schedule viewer and simulate the design in ModelSim by clicking the *SW/HW Co-Simulation* icon on the toolbar. You should see results similar to this:

```

...
# Cycle latency:      1040413
# ** Note: $finish    : ../simulation/cosim_tb.sv(287)
#   Time: 20808290 ns Iteration: 1 Instance: /cosim_tb
# End time: 16:25:24 on Dec 08,2020, Elapsed time: 0:00:16
# Errors: 0, Warnings: 0
...
Info: Verifying RTL simulation
...
Retrieving hardware outputs from RTL simulation for sobel_filter
function call 1.
PASS!
...
Number of calls:      1
Cycle latency:      1040413
SW/HW co-simulation: PASS

```

Observe that loop pipelining has dramatically improved the cycle latency for the design, reducing it from 3,392,549 cycles to 1,040,413 cycles in total.

Finally, use Microsemi's Libero to map the design onto the PolarFire FPGA by clicking the *Synthesize Hardware to FPGA* icon on the toolbar. Once the synthesis run finishes, examine the FPGA speed (FMax) and area data from the summary.results.rpt. You should see results similar to the following:

===== 2. Timing Result =====

Clock Domain	Target Period	Target Fmax	Worst Slack	Period	Fmax
clk	10.000 ns	100.000 MHz	7.275 ns	2.725 ns	366.972 MHz

===== 3. Resource Usage =====

Resource Type	Used	Total	Percentage
4LUT	1056	108600	0.97
DFF	607	108600	0.56
I/O Register	0	852	0.00
User I/O	0	284	0.00
uSRAM	0	1008	0.00
LSRAM	0	352	0.00
Math	0	336	0.00

5 Designing Streaming/Dataflow Hardware

The final hardware implementation we will cover in this tutorial is called a *streaming* implementation (also sometimes called a *dataflow* implementation). Streaming hardware can accept new inputs at a regular initiation interval (II), for example, every cycle. This bears some similarity to the loop pipelining part of the tutorial you completed above. While one set of inputs is being processed by the streaming hardware, new inputs can continue to be injected into the hardware every II cycles.

For example, a streaming hardware module might have a *latency* of 10 clock cycles and an II of 1 cycle. This would mean that, for a given set of inputs, the hardware takes 10 clock cycles to complete its work; however, the hardware can continue to receive new inputs every single cycle. Streaming hardware is thus very similar to a pipelined processor, where multiple different instructions are in flight at once, at intermediate stages of the pipeline. The word “streaming” is used because the generated hardware operates on a continuous stream of input data and produces a stream of output data. Image, audio and video processing are all examples of streaming applications.

In this part of the tutorial, we will synthesize a circuit that accepts a new input pixel of an image every cycle (the input stream) and produces a pixel of the output image every cycle (the output stream). Given this desired behavior, an approach that may spring to your mind is as follows: 1) Read in the entire input image, pixel by pixel. 2) Once the input image is stored, begin computing the Sobel-filtered output image. 3) Output the filtered image, pixel by pixel. While this approach is certainly possible, it suffers from several weaknesses. First, if the input image is 512x512 pixels, then it would take 262,144 cycles to input an image, pixel by pixel. This represents a significant wait before seeing any output. Second, we would need to store the entire input image in memory. Assuming 8-bit pixel values, this would require 262KB of memory. An alternative widely used approach to streaming image processing is to use *line buffers*.

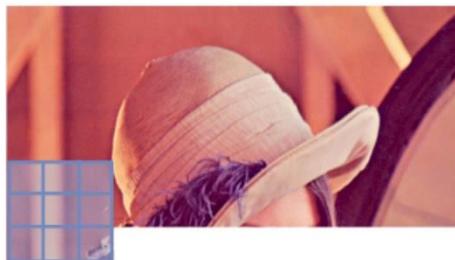


Figure 20: Motivation for use of line buffers.

Figure 20 shows the 3x3 Sobel filter sweeping across an input image. From this figure, we can make a key observation, namely, that to apply the Sobel filter, we do not need the *entire* input image. Rather, we only need to store the previous two rows of the input image, along with a few pixels from the current row being received (bottom row of pixels in the figure). Leveraging this observation, we can drastically reduce the amount of memory required to just two rows of the input image. The memory used to store the two rows are called “line buffers” and they can be efficiently implemented as block RAMs on the FPGA.

Create a new LegUp project for part 3 of the tutorial and include all the .cpp and .h files for part 3. Again, specify Microsemi’s PolarFire device (MPF100T-FCVG484I) and finish creating the project. Examine the sobel.cpp file in the project viewer and you will find the following line:

```
static LineBuffer<unsigned char, WIDTH, 3> line_buffer;
```

This statement “instantiates” LegUp’s LineBuffer template class from the <legup/image_processing.hpp> C++ library to create a line_buffer object. The template parameters specify the desired line buffer configuration: 1) use “unsigned char” 8-bit type to represent pixels, 2) set image width to WIDTH, and 3) set the filter size to 3. Inside the LineBuffer, there are internal arrays for storing the previous rows (two rows when filter size is 3), and an externally accessible 2D array named “window” to contain pixels in the current 3x3 receptive field. The line_buffer is declared as static so that its internal state and memory is retained between functions calls.

A few lines below you should see “line_buffer.ShiftInPixel(input_pixel);”. Each call of the ShiftInPixel() function pushes in a new pixel into the line buffer and updates the line buffer’s internal previous-row arrays as well as the receptive field window. In the subsequent nested loop, you will see the 3x3 receptive field is accessed by reading the “window” array, i.e., line_buffer.window[m + 1][n + 1].

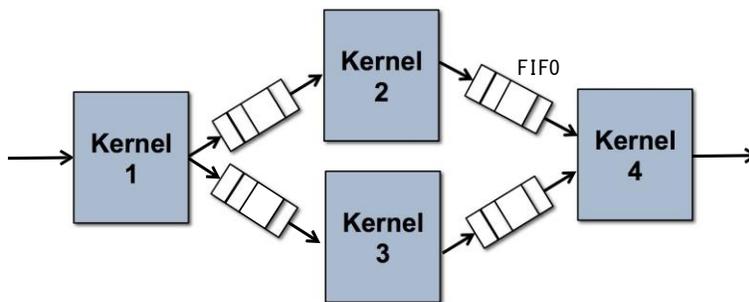


Figure 21: Streaming hardware circuit with FIFO queues between components.

Before going further, we will explain a common feature of streaming hardware called a FIFO (first-in first-out) queue. We use FIFO queues to interconnect the various streaming components, as shown in Figure 21. Here, we see a system with four streaming hardware modules, which are often called *kernels* (not to be confused with the convolutional kernels used

in the Sobel filter!). The hardware kernels are connected with FIFO queues in between them. A kernel consumes data from its input FIFO queue(s) and pushes computed data into its output queue(s). If its input queue is empty, the kernel stalls (stops executing). Likewise, if the output queues are full, the unit stalls. In the example in Figure 21, kernel 4 has two queues on its input, and consequently, kernel 4 commences once a data item is available in both of the queues.

The LegUp tool provides an easy-to-use FIFO data structure to interconnect streaming kernels, which is automatically converted into a hardware FIFO during circuit synthesis. Below is a snippet from the `sobel_filter` function in the `sobel.cpp` file. Observe that the input and output FIFOs are passed by reference to the function. A pixel value is read from the input FIFO via the `read()` function; later, a pixel is written to the output FIFO through the `write()` function. These functions are declared in the `legup/streaming.hpp` header file.

```
void sobel_filter(FIFO<unsigned char> &input_fifo,
                 FIFO<unsigned char> &output_fifo) {
    ...
    unsigned char input_pixel = input_fifo.read();
    ...
    output_fifo.write(outofbounds ? 0 : sum);
    ...
}
```

The rest of the `sobel_filter` function is very similar to the previous parts of this tutorial. An exception relates to the use of static variables so that data can be retained across calls to the function. A count variable tracks the number of times the function has been invoked and this is used to determine if the line buffers have been filled with data. Two static variables, `i` and `j` keep track of the row and column of the current input pixel being streamed into the function; this tracking allows the function to determine whether the pixel is out of bounds for the convolution operation (i.e. on the edge of the image).

In the main function in `sobel.cpp`, you will see that FIFOs are declared in the beginning. The FIFO class has a template parameter to specify the data type stored inside the FIFO. The FIFO constructor argument specifies the depth (how many elements can be stored). In this case, the FIFOs are declared to have the unsigned char data type to create 8-bit wide FIFOs.

In the main function, we see that the image input data (stored in `input.h`) is pushed into the `input_fifo` and the Sobel filter is invoked for `HEIGHT x WIDTH` times. Finally, the output values are checked for correctness and PASS or FAIL is reported. The main function returns 0 if the output values are correct. Click the icons to compile and run the software, and you should see the computed and golden pixel values and the message “*RESULT: PASS*”.

The `sobel_filter` top-level function has an additional pragma, “`#pragma LEGUP function pipeline`”, to tell LegUp that the `sobel_filter` function is intended to be a streaming kernel.

Now generate the hardware with LegUp by clicking the *Compile Software to Hardware* icon. In the report file (summary.legup.rpt) that opens up, you should see the top-level RTL interface now includes an input AXI stream interface and an output AXI stream interface, corresponding to the input_fifo and output_fifo arguments of the top-level function. Under Pipeline Result that the sobel_filter function is pipelined and has an initiation interval of 1.

===== 3. Pipeline Result =====

Label	Function	Basic Block	Location in Source Code	Initiation Interval	Pipeline Length	Iteration Count	Latency
sobel_filter	sobel_filter(legup::FIFO...)	%init.check	line 12 of sobel.cpp	1	7	n/a	n/a

This circuit has memories inside the hardware (see Local Memories under Memory Usage) due to the line buffers and the counters that are used. You can see there are two RAMs in the circuit both with 4096 bits, corresponding to the two line buffers each storing 512 x 8-bit pixels.

Local Memories					
Name	Accessing Function(s)	Type	Size [Bits]	Data Width	Depth
...					
sobel_filter_line_buffer_2_a0_a0	sobel_filter	RAM	4096	8	512
sobel_filter_line_buffer_2_a1_a0	sobel_filter	RAM	4096	8	512

Now Simulate the streaming hardware by clicking the *SW/HW Co-Simulation* icon. You will see scrolling output, reporting the computed and expected pixel value at each clock cycle. During this run, you should see messages like this:

```

...
MATCH: i = 511 j = 509 sw = 0 hw = 0
MATCH: i = 511 j = 510 sw = 0 hw = 0
MATCH: i = 511 j = 511 sw = 0 hw = 0
Result: 262144
RESULT: PASS
make[1]: Leaving directory '/cygdrive/c/LegUp-9.0/workspace/sobel_part3'
make[1]: Entering directory '/cygdrive/c/LegUp-9.0/workspace/sobel_part3'
Number of calls:      262658
Cycle latency:      262667
SW/HW co-simulation: PASS

```

The total number of clock cycles is about 262,667, which is very close to 512 x 512 = 262,144. That is, the number of cycles for the streaming hardware is close to the total number of pixels computed, which confirms that we are processing 1 pixel every clock cycle (II=1). At the end of the co-simulation, you should see that the co-simulation has passed.

Now, we can synthesize the circuit with Libero targeting the PolarFire FPGA by clicking the *Synthesize Hardware to FPGA* icon in the toolbar. You should see the following results in the summary.results.rpt file:

===== 2. Timing Result =====

Clock Domain	Target Period	Target Fmax	Worst Slack	Period	Fmax
clk	10.000 ns	100.000 MHz	6.003 ns	3.997 ns	250.188 MHz

===== 3. Resource Usage =====

Resource Type	Used	Total	Percentage
4LUT	605	108600	0.56
DFF	589	108600	0.54
I/O Register	0	852	0.00
User I/O	0	284	0.00
uSRAM	0	1008	0.00
LSRAM	2	352	0.57
Math	0	336	0.00

LegUp also allows the user to give a target clock period constraint, which the compiler uses to schedule the operations in the program and insert registers so that the generated circuit can be implemented accordingly. It may not always be possible for LegUp meet the user-provided target period precisely, due to the complexity of the circuit or the physical properties of the target FPGA device, but in general, a lower clock period constraint leads to higher Fmax. A lower clock period may cause larger circuit area due to LegUp inserting more registers, and a higher clock period constraint leads to lower Fmax but can also have less area.

The clock period constraint can be set in the HLS Constraints dialog as shown in Figure 22, select “Set target clock period” for *Constraint Type* and set *Constraint Value* to the desired clock period in nanoseconds.

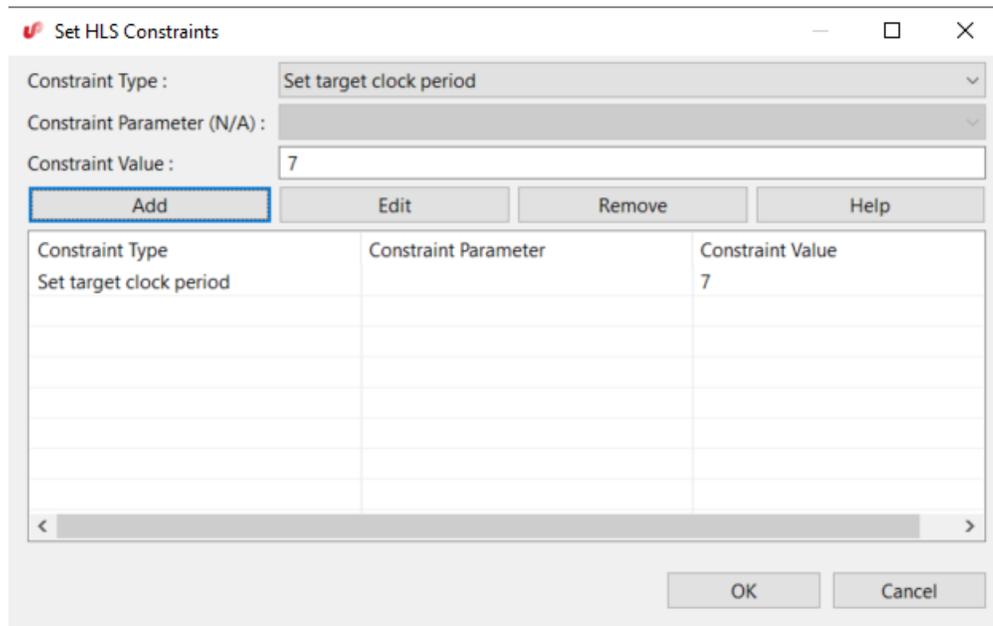


Figure 22: Setting the target clock period HLS constraint.

If the target clock period constraint is not provided by the user, as in this tutorial, LegUp will use the default target clock period constraint that has been set for each target FPGA device. The default clock period constraint is 10 ns for the Microsemi PolarFire FPGA.

Now you can try lowering the clock period constraint to 7 ns and recompile software to hardware. You should see the pipeline length has increased from 7 to 12 in summary.legup.rpt:

==== 3. Pipeline Result ====

Label	Function	Basic Block	Location in Source Code	Initiation Interval	Pipeline Length	Iteration Count	Latency
sobel_filter	sobel_filter(legup::FIFO...)	%init.check	line 12 of sobel.cpp	1	12	n/a	n/a

This is because LegUp has added additional pipeline stages/registers to achieve the higher target Fmax. You can also synthesize the generated circuit with Libero again to examine the impact of the clock period constraint on the generated circuit.

6 Instantiating LegUp IP Core in Libero SmartDesign

After we use LegUp HLS to design a hardware IP component, we will want to instantiate the component into Libero SmartDesign and integrate this core into our larger system. When LegUp HLS generates the hardware, LegUp will also generate a `create_hdl_plus.tcl` script to easily instantiate the LegUp-generated IP core into Libero SmartDesign. You will see the log message in the console which includes the full path to the script:

```
Info: Generating HDL+ Tcl script to be imported in SmartDesign:  
/cygdrive/c/LegUp-9.0/workspace/sobel_part3/create_hdl_plus.tcl.
```

This is a Cygwin path, you can replace `"/cygdrive/c/"` with `"C:\"` so the Windows path is:

```
C:\LegUp-9.0\workspace\sobel_part3\create_hdl_plus.tcl.
```

Now we can open Libero and create a new SmartDesign (or open an existing SmartDesign) as shown in Figure 23 below.



Figure 23: Create SmartDesign in Libero

We now go to the Libero *Projects* menu and select *Execute Script* and give the path to the generated `create_hdl_plus.tcl` script as shown in Figure 24. Then click *Run*. Running the Tcl script will add the LegUp-generated HDL+ component `sobel_filter_top` and all required Verilog files, memory initialization files, and other dependencies to the Libero project.

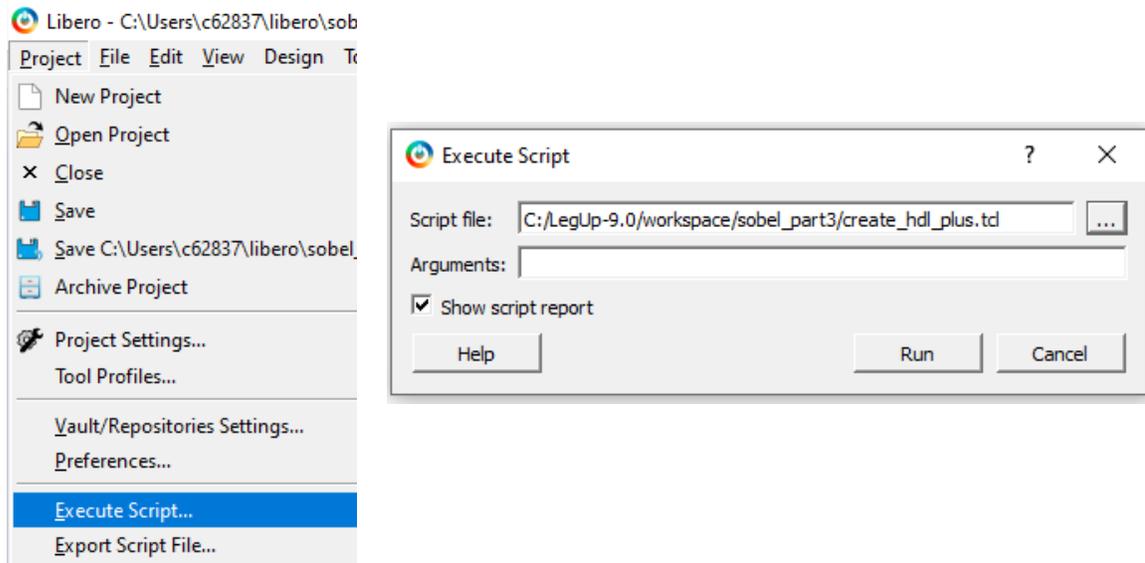


Figure 24: Execute Tcl script to instantiate LegUp-generate IP component into SmartDesign

As shown in Figure 24, we can now instantiate the component in SmartDesign by right-clicking on the `sobel_filter_top` HDL+ component in the *Design Hierarchy* panel on the left and selecting *Instantiate in system*. In the SmartDesign system we will now see the new `sobel_filter_top_0` IP component.

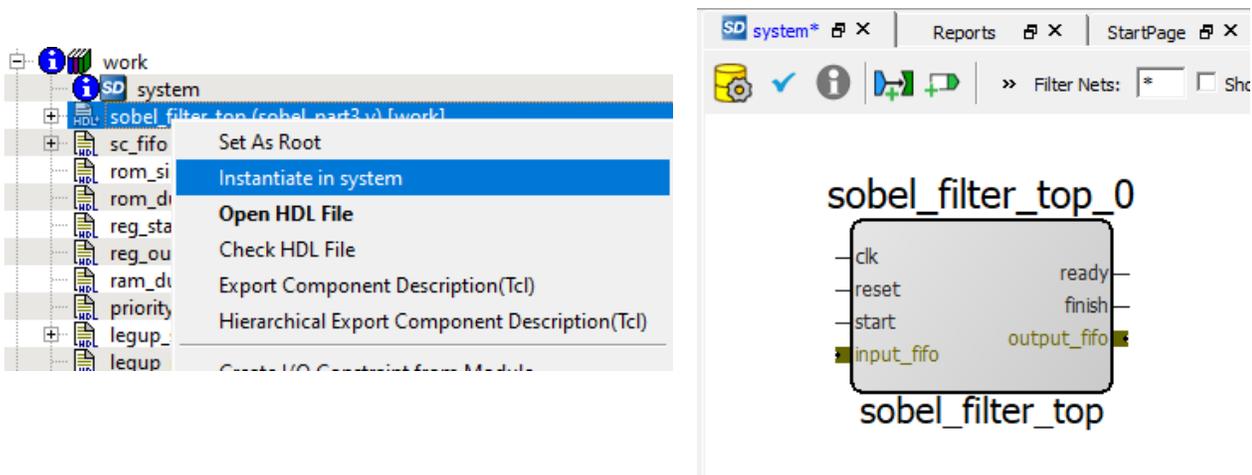


Figure 25: LegUp IP Component instantiated inside SmartDesign

In this case, since the `sobel_filter_top` IP component used LegUp FIFOs as top-level arguments, LegUp has automatically grouped the `output_fifo` and `input_fifo` data/ready/valid ports as AXI4-Stream interfaces.

7 Summary

High-level synthesis allows hardware to be designed at a higher level of abstraction, lowering design time and cost. In this tutorial, you have gained experience with several key high-level synthesis concepts in LegUp, including loop pipelining and streaming/function pipelining, as applied to a practical example: edge detection in images. These key techniques can allow you to create a high-performance circuit from software.

For any questions, please contact us at legup@microchip.com.