Introduction:

Convolution is a common operation in digital signal processing. In this project, you will be creating a custom circuit implemented on the Zedboard that exploits a significant amount of parallelism to improve performance compared to a microprocessor.

Although it is outside of the scope of the project instructions to give a detailed description of convolution (use google), I have included the pseudocode below:

```
for (i=0; i < outputSize; i++) {
    y[i] = 0;
    for (j=0; j < kernelSize; j++) {
      y[i] += x[i - j] * h[j];
    }
}
```

Convolution takes as input a signal (shown as the x array) and a kernel (shown as the h array). The output is another signal (y array), where each element of the output signal is the sum of the products formed by multiplying all the elements of the kernel with appropriate elements of the input signal. Note that this pseudocode is not complete. Specifically, the x array will be accessed outside of its bounds both at the beginning of execution where i-j is negative and towards the end of execution, where i-j is larger than the x array (the output size is sum of the input size and the kernel minus 1). Also, you will be using 16-bit unsigned integer operations, which need to be "clipped" to the maximum possible 16-bit value in case of overflow. See the provided C++ code for a correct software implementation.

The FPGA implementation will store the input signal in DRAM (currently a fake DRAM model implemented in blockRAM), and will read in a kernel (up to 128 elements) through the memory map. The FPGA will then execute like previous labs, using a go and size input from the memory map, while writing all results to another DRAM. Your datapath will fully unroll the inner loop, and then pipeline the outer loop. Due to resource limitations of the FPGA, the kernel size will be limited to 128 elements.

EDGE INSTRUCTIONS (All non-EDGE students must do the entire project):

EDGE students are allowed to use my provided solution to part 1, in addition to any other code that I provide. All other students must complete the entire project.

Part 1 – DRAM DMA Interface (EDGE students can use my solution, or can do this for extra credit)

For this project, you will be using the external DRAM on the Zedboard. However, because the current board template doesn't support external memory accesses yet, I have instead made a fake DRAM (dram_model.vhd) that is implemented on block RAMs, but exhibits many of the characteristics of DRAM.

The main challenge of interfacing with the DRAMs is dealing with both control and data signals that cross clock domains. The DRAMs run on a 133 MHz clock (**FCLK1**), while the rest of your circuit will run at 100 Mhz (**FCLK0**).

External memories are generally accessed through a DMA interface where the user specifies the starting address and the size of the transfer. In this part of the project you will be creating a DMA entity for reading from DRAM into your convolution pipeline (dram_rd_ram0). I have provided pre-synthesized solutions for the other dma entities (dram_wr_ram0, dram_rd_ram1, dram_wr_ram1).

For the dram_rd_ram0 entity, you will need to create an address generator for the DRAM. This will look a lot like your block RAM address generator from earlier labs. Note that the address generator should run on the same clock as the DRAM (133 MHz, FCLK1).

Next, you will need to integrate the DMA interface with the address generators. At the very least, you will likely have a go signal, a done signal, a starting address signal, and a size signal between the interface and address generator. Because the DMA interface is in a different clock domain (100 MHz, FCLK0), *all of these signals must be synchronized*. I strongly recommend using a handshake synchronizer between the interface and address generator to ensure that all the signals are stable before the address generator uses them. This is critical for reliable functionality. Without synchronization, the DRAM might occasionally work, but there will be non-deterministic behavior.

In addition to the previous signals, you will obviously need to also handle transferring data from memory into the DMA interface's data output that will be used by the pipeline. To implement the transfer of data across domains, use a FIFO with different read and write clocks. Also, the DRAM delivers 32 bits at a time, but your application works on 16-bit data. Therefore, your FIFO for dram_rd_ram0 should have a write width of 32 bits and a read width of 16 bits. Note that there are other ways of implementing this data transfer, but I would strongly suggest this one due to the buffering techniques that will be discussed later. For the input FIFO, make sure to use a programmable full flag that leaves at least 16 entries for outstanding read requests. Also, I'd recommend using the "first-word fall through" setting to simply the control logic.

Interfacing with DRAMs is not trivial, so I have provided a sample program that you can use to test your design. The provided code uses a datapath that simply passes data through unchanged from ram0 to ram1, which effectively allows you to test transferring data from one DRAM to another through the FPGA.

In wrapper.vhd, you will notice that there are actually 4 DMA entities. Two of these (dram_wr_ram0, dram_rd_ram1) are only needed because there actually isn't a DRAM, so we have to transfer data into it through the memory map. A 3rd entity, dram_wr_ram1, is provided to you due to time constraints.

To help you get started, I have provided presynthesized versions of all the DMA entities (see the *..edn files in the dram_cores directory). You use these cores by including it in your project and instantiating it in a VHDL port map, which is already done for you. Unfortunately, you can't simulate edn files, so I have also included a simulation model dram_rd_ram0_0_funcsim.vhdl. Similarly, there are also simulation models for the other 3 DMA entities. Note that the simulation model uses an extra _0 at the end of the entity name (dram_rd_ram0_0), so you will have to modify the component and port map in wrapper.vhd to use the simulation model, which you should then change back for synthesis.

I have provided a very basic testbench (wrapper_tb.vhd) to assist you, but be aware that you should make significant changes for testing purposes.

IMPORTANT: When instantiating the provided accelerator IP, there is a good chance that Vivado will have magically removed some of the critical files from the core. If it does not work, make sure to follow the procedure shown below:

UPDATE: Here is the procedure I used to get the project to compile with the accelerator.

1) **Put provided accelerator core in a repository directory.**
2) **Create a Vivado project like previous labs and instantiate both the Zynq and the accelerator.**
3) **Connect the clocks between the Zynq and the accelerator. The provided software requires you to connect FCLK0 to FCLK0, FCLK1 to FCLK1, etc. Note that this is different than lab 5 part 1.**
4) **Finish the block diagram with the automation options.**
5) **Edit the accelerator in the IP packager. For some reason the group files option will need to be updated, so select the merge changes option. Repackage the IP. Ignore the warnings about "File group does not contain any files."**
6) **Upgrade the IP when prompted.**
7) **Generate output products.**
8) **Create the HDL wrapper.**
9) **Generate the bitfile.**
10) **Upload to the board and test with the provided C++ code. If it works, the output should look like this:**

```
Starting job "./zed_app design_1_wrapper.bit" on board 192.168.1.107:
Programming FPGA....SUCCESS
Testing transfers to/from address 0....SUCCESS
Testing max transfer size....SUCCESS
Testing random sizes and addresses....SUCCESS
```

To simulate, you need to add wrapper_tb.vhd as a simulation source (after selecting edit in IP packager for the accelerator). Similarly, add the four *_funcsim.vhdl files in src/dram_cores as simulation sources. Modify wrapper.vhd by adding a _0 to the end of the 4 dram_rd/wr components and port maps. As mentioned before, the simulation models have the _0 suffix, so make sure you are using the appropriate names depending on whether or not you are doing simulation or synthesis. Next, simulate the provided testbench and see if there were any errors. Note that the provided testbench does a very basic test. You should expand it significantly if your code isn't working on the board.

There are a lot of things that can go wrong here, so start as soon as possible so I can help.

Part 2 – Convolution pipeline

Due to time constraints, I will be providing a generic solution to everyone (see mult_add_tree.vhd). However, you are still welcome to do this as extra credit. If you use my code, make sure to instantiate the unsigned_arch architecture.

The pipeline for convolution is large, but conceptually simple. In this project, we will limit the maximum kernel size to 128. To support this kernel size, the first row of the pipeline consists of 128 16-bit multipliers, each of which will multiply corresponding elements from the signal and kernel. After the first row, an adder tree is used to add all the products together. The output of the adder tree (i.e., final sum) defines a single output element, which is the output of the datapath.

If you decide to implement the pipeline yourself, there are several things to be aware of. First, the pipeline should have a register after every multiplier or adder. Second, the datapath takes 16-bit inputs and produces 16-bit outputs, which means you don't need to increase the width of operations as you go deeper into the datapath. However, you will need to implement saturation (e.g., clipping of an audio signal). In other words, if at any point in the datapath the result of an operation exceeds 16 bits, the output of the operation should be all 1's (0xffff). If you don't implement this clipping, the results will wrap back around to zero, which will not be correct. There are two ways to do this: 1) implement saturation within each adder/multiplier, or 2) perform the multiply-add tree while storing all overflow bits, and then implement saturation on the pipeline out. Unless you want to modify the code I provided, I recommend the 2nd method.

The entire pipeline consists of 128 multipliers, 127 adders, and a lot of pipeline registers. If you implement your own pipeline, I highly recommend using VHDL generate statements and arrays, otherwise your code will be huge. You can actually use recursion to create very concise structural architectures, which is demonstrated by my provided code.

Part 3 –Signal Buffer

You might have noticed that there is potential problem in the first two parts of the project. The DRAM only delivers 32 bits at a time, and the DMA interface only outputs 16 bits at a time, but the pipeline needs 128*16 bits every cycle to avoid stalls. Fortunately, there is significant overlap between iterations. In fact, convolution is a sliding-window algorithm where the window moves by one element each iteration. Therefore, we can exploit the overlap between iterations to reuse data and improve bandwidth. There will be a class lecture soon that discusses the exact details.

To enable data reuse, you will need to create a simple buffer that generates 128-element signal windows for the pipeline. Because the window only differs by a single element each iteration, the buffer can be implemented as a large shift register that shifts in 16 bits at a time.

You are free to implement the buffer however you like, but I'd recommend using a FIFO-like interface that specifies if the buffer is empty/full. With this interface, the pipeline can read whenever the buffer isn't empty and you can write data into the buffer (from the ram0 read DMA interface) whenever the buffer isn't full. The data input to the FIFO should be 16 bits, and the output should be an array of 128 16-bit elements. The easiest way of handling this output is by creating an array type that is defined in a package (e.g., user_pkg.vhd), so that you can use the array in the port definition. Assuming you implement everything correctly, this buffer will be capable of delivering a 128-element window to the pipeline every cycle (assuming data is available from the DRAM). If your speedup is significantly less than what is demonstrated in class, then the problem is likely that this buffer does not provide a new window every cycle,

Part 4 –Kernel Buffer

In addition to buffering the signal as it is read from memory, you need some way of storing and accessing the entire kernel. Although we could potentially read the kernel from memory, because I have kept things simple by limiting the kernel size to 128 elements, you can transfer the kernel into registers in the FPGA using the memory map. In the provided memory map for convolution, there is an address, KERNEL_DATA_ADDR, that is used for these transfers. You should use another shift register that consists of 128 16-bit elements, which shifts in 16 bits every time that the memory map writes data to an address that corresponds to the kernel. After 128 memory map transfers, the entire kernel should be loaded. Note that this functionality is identical to part 3, so ideally you can reuse the same buffer entity. It would probably be a good idea to check if the kernel is actually loaded during testing. To support this, I included a

KERNEL_LOADED_ADDR in the memory map that you can use to read the status of the kernel buffer. This is completely options, but provides information for debugging if you want to use it.

Other tasks

In addition to the previously discussed components, you will also need a simple controller, glue logic, etc. I have provided the memory map (memory_map.vhd in the convolve directory), but you might need to modify it based on your implementation of other components. I would highly recommend basing your design on the provided dram_test code. Notice I have removed the starting addresses in the convolution memory map because we can assume that the signal and output both start at address 0. Make sure you use the provided memory map in the convolve directory because the dram_test memory map is not compatible with the convolve software.

After you have implemented all components, modify the provided testbench to test your code. **Note that the provided testbench is intended to just get you started. It is up to you to extend it for thorough testing. It currently provides no automatic checking for errors.** When you are confident that it is working, create a bitfile. Test your implementation with the provided C++ code on the class server. You cannot modify the C++ code to fix problems in the VHDL, but you are welcome to change it for debugging and testing purposes.

The C++ code does a variety of tests ranging from simple tests that are likely to work up to stress tests that use the largest possible signals with random values. If your code works properly, you should see an output like this:

```
Starting job "./zed_app design_1_wrapper.bit" on board 192.168.1.102:
Programming FPGA....Testing small signal/kernel with all 0s...
Percent correct = 100
Speedup = 0.0159574

Testing small signal/kernel with all 1s...
Percent correct = 100
Speedup = 0.0205479

Testing small signal/kernel with random values (no clipping)...
Percent correct = 100
Speedup = 0.0205479

Testing medium signal/kernel with random values (no clipping)...
Percent correct = 100
Speedup = 3.07368

Testing big signal/kernel with random values (no clipping)...
Percent correct = 100
Speedup = 15.9339

Testing small signal/kernel with random values...
Percent correct = 100
Speedup = 0.0196078

Testing medium signal/kernel with random values...
Percent correct = 100
Speedup = 2.80223
```

```
Testing big signal/kernel with random values...
Percent correct = 100
Speedup = 14.5488

TOTAL SCORE = 100 out of 100
```

Extra Credit:

If you finish early, there are quite a few options for extra credit. In the current implementation, you should notice that all padding of the signal (i.e., handling the cases where the bounds of the signal are exceeded) and kernel (i.e., handling kernels less than 128 elements) is performed in software. In other words, 0's are added to the beginning and end of the signal, and kernels less than 128 elements have 0's added to the beginning. Padding in software works, but is potentially slow, especially for large signals. One good extension for extra credit is to implement the padding in the FPGA itself. Note that this will also require changes to the software code.

If you really want to impress me, you can also extend the FPGA implementation to handle arbitrary kernel sizes. I'll warn you that this is not trivial, and I won't give you the exact techniques required, but here are a few hints. To implement an arbitrarily sized kernel, you must perform the convolution multiple times, each time with a different 128-element segment of the kernel. The tricky part is making sure the padding of the signal is done correctly each time. In addition, the output of each iteration is a partial result, which must be accumulated with multiple future iterations to get the correct answer. Do not attempt this step until you have finished everything else.

Another possibility would be to convert your implementation to use floating point instead of integer operations. You will not be able to fit as many operations on the device, but if you implemented your VHDL in a good way, this should not require many changes to your code.

You could also create a frequency-domain implementation that uses an FFT and IFFT. I won't explain this implementation, but there is plenty of information on google.

I would also like to see the code demonstrated on a real example. The use of 16-bit signals and kernels makes this project perfect for testing 16-bit audio. Note that you might need to make some changes to the pipeline to support different types.

All extra credit should be turned in using a separate directory so there is no confusion when grading the original project.

Report Instructions: IMPORTANT

If your project is not 100% complete, it is up to you to show me what you have completed. Therefore, you should include a report that shows testbench simulations, explanations, etc. that demonstrate that certain parts of the project are correct. Do not simply say that "we finished the DRAM interface", or you will not get credit. However, if you show detailed waveform simulations and correct output from the provided test application, then I will be convinced. If you have fully completed the project, the report can simply be a brief description of your implementation, a summary of any problems you encountered, and how you fixed those problems. If you did any extra credit, it should be described in the report.

SUBMISSION INSTRUCTIONS (One submission per group)

Make sure the names of all group members are at the top of every file that you create and/or modify.

Create a directory with the name of your server account. Use the following structure:

```
server_account_name/
        readme.txt              // Group members, anything that the grader needs to
                                // be aware of
        report.doc or pdf       // detailed report explaining what works (with
                                // convincing evidence)
        convolve.bit
        accelerator_1.0/        // IP core from repository with this exact name
                                // Make sure all VHDL is included
        extra_credit/           // If applicable, put any extra credit files in here.
                                // Include a separate report in this directory that
                                // explains the extra credit.
```

Zip the entire directory and submit the zip file. Note that you do not have to submit any C++ code, unless you did something for extra credit.