# LAB 8 – Full System Integration

Goals

- Learn how a processor is built.

- Learn how the processor communicates with the outside world.

- Implement the MIPS processor and demonstrate a simple "snake" program on the FPGA starter kit.

To Do

- Learn the components of the MIPS processor. The ALU was designed in Lab 5. Several other components are given.

- Assemble the components of the MIPS processor.

- Make the changes to the processor so that memory mapped I/O write operations can be performed.

- If everything is correct, you should observe a tiny "snake" that loops on the 7-segment LED.

- Make changes to the system to control the speed at which the snake crawls. The speed should be based on the amount specified by the switches on the board.

- *[Optional] For the challenge seekers, try to change the snake's motion pattern. Additionally, you could try to change the direction of snake's motion based on a switch input.*

- Follow the instructions. Paragraphs that have a gray background like the current paragraph denote descriptions that require you to do something.

- This lab spans **two lab sessions**. In total, you can get up to **6 points plus 2 bonus points**.

- To complete each lab session, you have to **show your work to an assistant before the deadline.** You will get **up to 3 points for each lab session.**

- You may additionally **hand in the Lab reports** that are provided for this lab. You will use **Moodle for the submission.** You will get **up to 1 bonus point per submitted Lab report. Only one submission per group is necessary.** The names of all members of the group must be in the report.

- As indicated in Lab, you have to **show Part 1 (Session I) to an assistant**. In addition, **you may submit Parts 1-4 in the Lab Report (Session I) for an additional bonus point via Moodle** (https://moodle-app2.let.ethz.ch/mod/assign/view.php?id=240199).

- As indicated in Lab, you have to **show Part 2 (Session II) to an assistant**. In addition, **you may submit Parts 1-3 in the Lab Report (Session II) for an additional bonus point via Moodle** (https://moodle-app2.let.ethz.ch/mod/assign/view.php?id=240200).

- All optional tasks are highly recommended. You can ask assistants for feedback on optional tasks.

Introduction

This is it. This is the exercise we have been working towards the entire semester. During this week's exercise, we will finally put together a microprocessor and run your own programs. In order to see the processor in live action, we will add some I/O interfaces to control and display a crawling snake on the 7-segment LED. Figure 1 shows a simplified block diagram of how the final system looks. This lab is divided into **two parts** spread over **two weeks**. In the first part, we put together the basic building blocks of a processor and try to run a "crawling snake" program. In the second part, we extend the I/O functionality to control the speed of the crawling snake using the switches present on the FPGA starter kit.
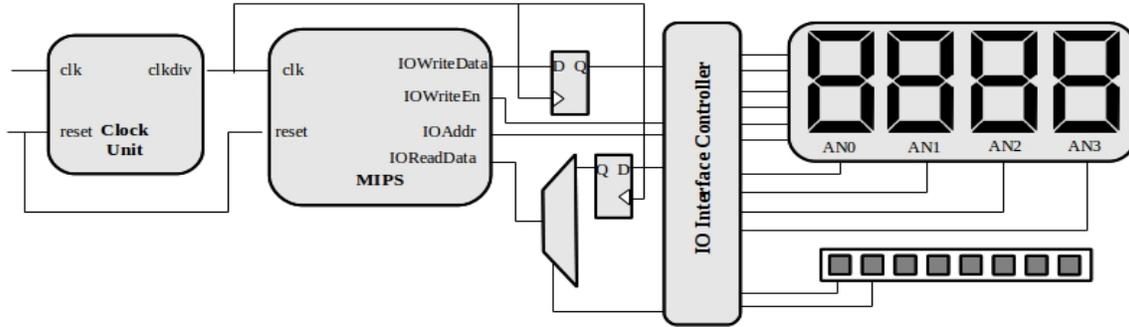
Figure 1. Simplified block diagram of this exercise

The MIPS processor that we will implement is closely based on the architecture presented in your textbook; specifically, the single-cycle processor depicted on page 383 in Figure 7.11. In the example code, we adhere to the same signal names and organization as the figure in the book to make it easier to follow the given source code. The processor itself is not terribly complicated, and its performance pales in comparison to a modern Intel processor. However, it is far from being useless. The MIPS that we implement runs faster than 10 MHz, and can execute 32-bit instructions. This easily beats state-of-the-art processors from the early 1980s (e.g., Intel 8086 and Motorola 68000). For example, anything that landed on the moon had a much less capable processor. The system clock on the FPGA board is 50 MHz, which is slightly faster than what we can use for our processor (the critical path is around 25 ns); but don't worry, the provided clock_div block generates a clock that is 4 times slower (12.5 MHz), so you do not have to design it separately.

While the processor itself provides interesting functions, it alone cannot communicate with the external world. It requires input and output interfaces. To this extent, we modify the MIPS processor to have standard I/O signals. In this standardized interface, we have a separate data output (`IOWriteData`) and data input (`IOReadData`) bus (each 32 bits wide). Since we want to address multiple I/O resources we will also provide an address signal `IOAddr`. This address helps us to identify which I/O resource we are accessing (writing to the 7-segment display or reading from the register that holds the direction input, etc.). Finally, there is a signal `IOWriteEn` that tells us that we are actually performing a write operation (not a read I/O operation).

We use two registers (the challenge requires another register): (i) to store the current value to be displayed on the 7-segment display and (ii) the speed level for the crawling snake. We design the I/O interface controller that ties all these circuits together and enables the processor to access the various I/O circuits that we have added.

It is important to note that we could do most of these changes also inside the processor. We add these changes into another hierarchy because we do not want to re-design the processor every time we have a new I/O device. In other words, the MIPS processor (with the I/O interface) stays constant even if we completely change the I/O circuits.

SESSION I

THE CRAWLING SNAKE

In this first part, we put together the building blocks of the MIPS processor. For this exercise, we provide you a Vivado project that already contains many parts of the processor. Once you examine the code, you realize that there is nothing very mysterious about any part of it, in fact you could easily write all of it without problems. However, the instructions will guide you through the exercise by explicitly saying which parts of the code are relevant for which task. Note that the given code and architecture are not necessarily perfect.

Go to the course webpage and download the .zip file containing the archive with the Vivado project directory. Extract the directory in your working directory and start Xilinx Vivado. Open the project file lab8.xpr that is among the extracted files.

In the directory of the extracted files, you will find the following file structure:

top.v	Top level hierarchy that connects the MIPS processor to the I/O on the FPGA board. *You will modify this file for Part 2.*
top.xdc	Constraints file of the top level. *You will modify this file for Part 2.*
MIPS.v	The main processor. *For Part 1 you have to add stuff inside of this file only.*
DataMemory.v (datamem_h.txt)	The initial content of the data memory (composed of 64 32-bit words). *The datamem_h.txt file contains the data part of the assembly program in a hexadecimal form. This module "loads" the data. You will only have to modify the .txt file if you do the challenges.*
InstructionMemory.v (insmem_h.txt)	The ROM (composed of 64 32-bit words) that contains the program. *The insmem_h.txt file contains the assembly instructions we want to run on the MIPS processor in a hexadecimal form. This module "loads" the instructions. You will modify the .txt file for Part 2.*
RegisterFile.v	Register file that creates 2 instances of *reg_half.v* as read ports and has 1 write port. *Don't worry about this one, it's just the implementation of a register.*
reg_half.v reg_half.ngc	Component describing a single port memory and binary description of how it is mapped in the FPGA. *Don't worry about these ones, they're used to implement the register.*
ALU.v	ALU similar to the one from Lab 5. *You should not change anything in this file, but if you want you can use your own implementation (just make sure that it works).*
ControlUnit.v	The unit that does the instruction decoding and generates nearly all the control signals. Table 7.5 on page 379 lists most of them and their truth tables (only the `AluOp` signal is generated differently in the exercise). *This is just a combinational circuit and it's already given, you don't need to change anything here.*
snake_patterns.asm	Assembly program corresponding to the *datamem_h.txt* and *insmem_h.txt* dump files that displays a crawling snake on the 7-segment display when all the parts are connected properly. *You have to modify this file for Part 2, where you'll also learn how to generate the dump files.*

MIPS Processor

In this part of the exercise we build the MIPS processor, which is implemented in the provided file *MIPS.v*. This file, however, is incomplete since the main blocks are not yet connected (instantiated). Figure 2 shows the corresponding the block-level overview of the MIPS processor. Note that the output signals are not connected because later in this exercise you have to decide how they should implemented according to their description.

To complete building the MIPS processor, you have to finish the following tasks.

Open the file *MIPS.v*. Note that all the required signals are already declared at the top of the module. Use the block diagram in Figure 2 as a reference to add the correct instantiation for:

- **Instruction Memory**: Note that the address of the instruction to be read is determined by the PC (program counter). The PC is always incremented by 4 to fetch the next instruction from the memory. We add 4 instead of 1 because each address in memory stores one byte and each MIPS instruction requires four bytes of memory. Therefore, we can throw away the 2 least significant bits of the address (because they are not necessary here) and use the 6 most significant bits (7 to 2) for its 64 words.
- **ALU**: The given (or your) ALU from Lab 5 has 4 bits for `AluOp`, whereas the controller generates a six-bit value that is the function field of an R-type instruction. You will have to select the 'correct' four bits to connect here[1].
- **Data Memory**: Just like the instruction memory, use the 6 most significant bits (7 to 2) of the actual address.
- **Control Unit.**

Memory Mapped I/O

In order to see the processor in action, we must extend the MIPS processor previously built to communicate with the external peripherals (e.g., 7-segment LEDs, switches, buttons). We will use a simple memory mapped I/O architecture, i.e., part of the memory address space is reserved for I/O operations. When data is written to or read from this address space, it is intercepted by the circuit and re-routed to the I/O circuitry. To complete this part of the exercise you must complete the signal assignments for `IsMemWrite`, `IOWriteData`, `IOAddr`, `IOWriteEn`.

Open the file *MIPS.v* and complete the assignments. The code is sufficiently commented to help you complete the incomplete signals. For example, the `IsIO` signal has already been implemented. In our processor, we reserve the address range 0x00007FF0 to 0x00007FFF for I/O operations.

---

[1] At the moment, we are only using 7 ALU instructions, and the ALU does not need all 6 signals of the Function field (of an R-type instruction). In Lab 9, we will need to modify the ALU a little bit and add support for more bits. We wouldn't want to also modify the Controller at that time, so we prepare for the future.

Fig 2. Block diagram of the MIPS processor that we will implement in this exercise. This diagram is almost identical to Fig 7.14 on page 387 of your textbook

5

Crawling snake on the 7-segment LED

In this first part of the exercise, we only need a 28-bit output register that contains the value to be sent to the four 7-segment LEDs. We assign a separate address for this register as given in the table below:

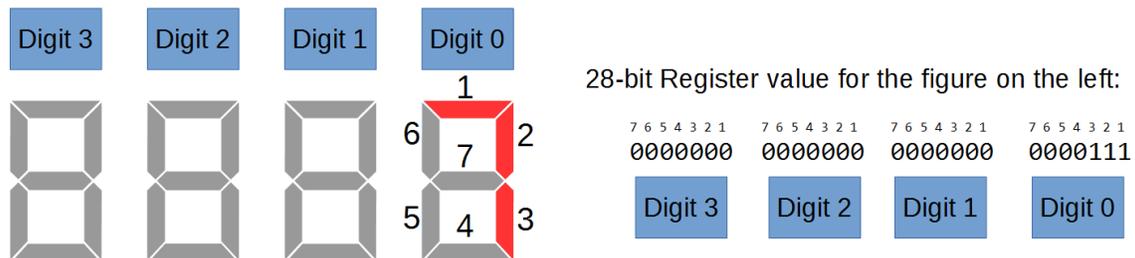| Address	Direction	Width	Description
0x00007FF0	out	28 bits	Value to be sent to the 7-segment display.

Therefore, the following instruction in assembly program:

```
sw $t0, 0x7FF0($0)
```

writes the content of the register $t0 to the address 0x00007FF0[2], which represents this other register outside of the processor. Additional external circuitry (already implemented in *top.v*) converts these values so that they can be displayed on the FPGA board. In this exercise, we will not use the decoder circuit from Lab 3, since there are no digits to be displayed. We directly use the contents of the 28-bit register to simulate the crawling snake.

It would be easy if the four 7-segment LEDs had individual signal lines. In practice, displays are rarely driven by such parallel connections. Since humans can only distinguish movement that is slower than 20-50 ms, each of the four 7-segment LEDs is enabled sequentially. On the FPGA board, all four displays have a separate enable pin (AN3, AN2, AN1 and AN0), and the segment connections are shared among all displays. The enable pins are active low, meaning that the segment displays the number if the corresponding AN signal is 0 and does not display anything if it is 1. To display all four numbers, we need four clock cycles. In the first cycle, AN3 is set to 0, all others are set to 1 and the pattern to be displayed is applied. In the next clock phase only AN2 is activated and the second number is applied, and so on. Once the last number is displayed the loop starts with the first number again. We will see our pattern as long as this process is repeated swiftly enough (faster than 20 ms). This logic has already been implemented for you in the *top.v* file and the *top.xdc* contains the correct mapping.

The content of the 28-bit register defines the 7-segment display as follows:

28-bit Register value for the figure on the left:

```
7 6 5 4 3 2 1   7 6 5 4 3 2 1   7 6 5 4 3 2 1   7 6 5 4 3 2 1
0000000         0000000         0000000         0000111
```

Digit 3   Digit 2   Digit 1   Digit 0

---

[2] We are using the 'Compact Data at Address 0' configuration of the MIPS. In this configuration, the MMIO address range is from 0x00007FF0 to 0x00007FFF. We will use this range as it is easier for us (only 16-bit addresses are used).

Assembly Program

We have a processor that is able to do memory mapped I/O; we have added an interface that uses the 7-segment display and several buttons. Now we need to make everything work. For this purpose, we need an assembly program that implements the crawling snake. This has already been implemented for you in the *snake_patterns.asm* file. Try to understand the code, it has been sufficiently commented. The files *datamem_h.txt* and *insmem_h.txt* contain the data and instruction op-codes corresponding to the file *snake_patterns.asm*. We will learn how to generate these hexadecimal dump files in Part 2.
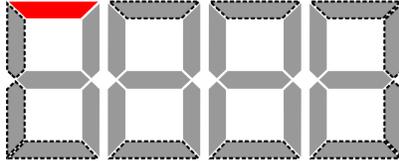
Generate the programming file in Xilinx Vivado and program the FPGA. You should see a crawling snake moving as shown in the figure above. Fix any problems you might have. Show your results to an assistant.

SESSION II

SPEED UP THE SNAKE

In this part of the exercise, we will try to control the snake's speed. The idea is to use two switches on the starter kit to control 4 different speed values. The tasks are the following:

1. Extend the top-level hierarchy to accept inputs from switches and transfer the values to registers.

2. Understand the provided assembly program *snake_patterns.asm* and modify it to accept a 2-bit input value from the switches. Based on the input, increase or decrease the speed of the snake.

3. *Optionally, you have two challenge tasks to complete.*

Extending the I/O functionalities

For the crawling snake exercise in the previous part, we only had to store the value to drive the 7-segment display appropriately. In this part, we need to read the values from the switches on the starter kit. For this purpose, we add an additional I/O register to the MIPS processor that is 2 bits wide (1 bit for each switch) and assign an address. The table of I/O registers is as follows:

Address	Direction	Width	Description
0x00007FF0	out	28 bits	Value to be sent to the 7-segment display.
0x00007FF4	in	2 bits	Speed step 0, 1, 2 or 3.

In the top.v file, create an input signal for the switches. Depending on the value of `IOAddr` set the `IOReadData` signal that connects to the MIPS processor. Modify the XDC file so that the input signal is connected to the correct FPGA pins. Synthesize the project to check for any errors. Note: `IOReadData` is a 32-bit signal, however we use only the 2 LSBs to read the status of the switches.

Pin	Label on board	Description
V16	SW1	Slide switch, MSB of speed step amount.
V17	SW0	Slide switch, LSB of speed step amount.

Modifying the Assembly program

The assembly code from Part 1 simply loops continuously with no controls. Therefore, we need to modify the code to accept the input from the switches and accordingly control the speed of the crawling snake.

Using the MARS environment, modify the assembly program *snake_patterns.asm* so that it will read in the 2-bit switch value using memory mapped I/O. Then, depending on the input, modify the delay loop such that the snake moves faster or slower.

Hint: You can use the MARS environment to check your modified code. Just remember to test it with a smaller loop counter value.

Your modified program needs to be copied into the text file *insmem_h.txt*. To do this, you can use the Memory Dump option within the MARS assembler. By selecting the "Memory Segment" as ".text", and "Dump Format" as "Hexadecimal Text", you are able to generate the required file. All you have to do is use a text editor and make sure that the file has exactly 64 lines. All lines after your real code have to be filled with zeroes.

Since we use memory-based variables in the data memory, the ".data" segment has to be exported to *datamem_h.txt* the same way as the ".text" segment. Notice that the *datamem_h.txt* file, after exporting, contains many lines. Make sure that the file has 64 lines, and delete all lines after that.

Generate the programming file in Xilinx Vivado, program the FPGA and control the speed of the crawling snake using the switches. Fix any problems you might have and show your results to an assistant.

Summary of the Design Flow

The following figure summarizes the steps to follow while building the system. The workflow can be divided into hardware and software sections. The software section consists of the assembly program that the MIPS processor executes. The hardware section comprises of the MIPS processor and the peripheral system itself. The assembly code is written and validated using the MARS simulator environment. After checking if the code produces the desired output, the data and instruction memories are dumped as hexadecimal values. The two memory dumps represent the data and opcodes that are used by the MIPS processor. The hardware architecture is constructed in Verilog using Xilinx Vivado. The *DataMemory.v* and *InstructionMemory.v* files use the files dumped by MARS to instantiate the memory blocks. The XDC file is modified to map the circuit to the FPGA board. Finally, the bit file used to program the FPGA is generated. Remember to follow the steps when you make any changes to the files.

The Challenge

We now have a complete system consisting of a MIPS processor that can run our own programs, accept inputs from the on-board switches, and display results on the 7-segment display. Here are some ideas for the challenge seekers.

Challenge 1: Change the snake crawling pattern. For example, modify the assembly code to crawl the snake in a zig-zag pattern as shown in the figure below.

Challenge 2: Toggle the direction of the snake. Add an additional switch input which will control the direction of the snake's movement. For example, the snake moves forward (you can choose to either loop or zig-zag) if the switch is 0 and backward when the switch is 1. This is more challenging because of the limited number of instructions that can be executed in our MIPS processor. You will also need to modify the top-level hierarchy to add a register to transfer the value of this new switch.

Last Words

In this exercise, we have realized a complete and working microprocessor that is able to execute a small subset of the MIPS instruction set. It doesn't seem to be that complex, does it? There are two main reasons why writing the code of the processor was (relatively) easy.

1. First, MIPS uses a simple architecture and one of its priorities is a simple design implementation.

2. Second, we started from a well-documented block diagram (taken from the book) that had already decided how the processor was partitioned, what signals were used to connect them, and how the signals were defined. All that was left was to convert the block diagram into Verilog syntax.

Additionally, we added the capability to interface with the rest of the world so that we can transfer the processor to the FPGA board and actually see our programs execute.

In the following lab, we will improve the current processor by adding new instructions and speeding up calculations.