

Patrick Murray

4/17/2012

EEL4713

Assignment 5

Introduction

For assignment 5 we were tasked with creating a fully functional pipelined MIPS processor. I used the code and general structure of assignment 4 to design the multi staged version in this lab. There are a number of benefits to using a pipelined datapath as opposed to a single cycle datapath, the most significant being a large speedup. We are able to produce a faster processor because we don't need to wait for every instruction to complete before beginning the next. In a single cycle processor, after we are done reading the data from a register and inputting it to the ALU we aren't doing anything else with the register until we need to write back to it. Using a pipelined version enables us to break the datapath up into five separate stages, allowing us to work on up to five instructions at once. This cuts out the slowdown of the critical path that we saw in our single cycle design.

Each one of the five stages mentioned above contains parts of the original single-cycle datapath with additional hardware needed to operate the pipelined version. The five stages consist of Instruction Fetch (IF), Instruction Decode (ID), Execute (EX), Memory (MEM), and Write Back (WB). Below is a short, general overview of each of the stages.

- IF
 - The main components in this stage are the Program Counter (PC), PC incrementer/adder, and the Instruction Memory. This stage receives signals from the MEM stage as well because that is where the branch mux is located, which supplies the PC with the next address to load.
- ID
 - This stage consists of the register file, controller, and jump mux. All the control signals are decoded and generated here. Because of this the register that is between ID and EX is very large as most of the control signals have to pass through here. This stage also contains the extending unit to either sign extend or zero extend immediate values. It receives data and signals from the WB stage because of the register file needing to be written back to.
- EX
 - This stage contains the ALU, forwarding muxes, and the JR mux. It also has a mux that takes care of the register destination which is routed through the entire pipeline and back to the ID stage to tell the register file what address to write to.
- MEM
 - This stage contains the branching logic, branch mux, data memory, and decoding unit. As mentioned previously, this is where the incremented PC signal departs from. The output of the branch mux will contain either the PC+4 signal, address we want to jump to, the address to return to after a JR, or the address to branch to.
- WB

- This stage is the smallest and only contains the mux to determine which output to route to the data of the register file in the ID stage. It essentially does exactly what its name implies, just writes all the data back to previous stages if it hasn't already been sent back.

In my datapath design I used large register files in between the five stages. Each one of these register files took in inputs from the previous stage and passed them through to the next stage. The signals that were passed through were either needed in a future stage or following the path back to the register file. I further broke these register files up into smaller files. The main register handled the unique signals going from stage to stage, while smaller register files took in the control signals only and carried them to their appropriate stages. You can see this in the sample datapath image at the end of this section.

While a pipelined design is much faster it also presents a number of new challenges. First the single cycle design needs to be converted to a pipelined datapath by adding large registers between our five cycles to stall the signals every clock cycle. This can become cumbersome because we need to pass certain signals through the registers, some signals need to avoid the registers, and some signals need to be passed backwards to over cycles from various later stages.

The main problem we need to deal with when designing the pipelined system is data hazards. There are many different hazards that need to be resolved and just figuring out which instructions create a hazard is a difficult task. A hazard can come in two general forms. The first is resolved using stalling. If we have an instruction that relies on data from another instruction which hasn't completed yet and needs to complete either completely or partially, we will need to add *nop* instructions to stall it. Stalls are mainly used when we use a load instruction (*lhu*, *lbu*, *lui*, *lw*) followed closely by another instruction that uses the data that is being loaded into our register file still. Stalling has the downside of slowing down our system. A *nop* is essentially just a blank instruction that does nothing, so inserting them into our datapath creates more instructions. The other hazard is resolved using forwarding. This hazard is again encountered when an instruction relies on a recent past instruction, but instead of needing to stall the entire pipeline and wasting those extra cycles, we grab the data we need from a future cycle and route it around the registers to the cycle that needs it, injecting it into the necessary signal path. This preserves our increased speedup but has the downside of being difficult to implement for the designer. I needed to again realize all the instruction combinations that would require forwarding, then build a forwarding unit into my datapath, route all the required signals into and out of it, add any additional blocks into the cycles or datapath to keep functionality, and design the logic to not only keep our initial design working but also the added forwarding and stalling.

An example of a hazard would be if the program added to values and stored the result into register 1. If the next instruction uses register 1 as a source register than the data will not yet be available to it. As I mentioned before there are a number of combinations of hazards, but primarily you need to look at the current instruction and the 3 instructions following it. If any of the 3 instructions that follow the current instruction use its value you are likely to run into a hazard. This can be slightly reduced depending on how you design your datapath. Instead of clocking my register file on the rising

edge, I instead used the falling edge. This allows you to not have to worry about any instructions that follow the current instruction by 3 spots (though you still need to take care of instruction 1 or 2 places ahead). This works by allowing the WB stage to output the data needed at the register file in time. The WB operates off of the normal rising edge clock and sends its data to the register file write port. Because the register file waits until the falling edge the data that it needs is already available and no forwarding is required.

To implement the forwarding in hardware you need a unit that handles the logic of stalls and forwards as well as a mux in front of both inputs to the ALU. Again this can vary depending on how the datapath is designed. The hazard table given to us shows that we actually need two muxes, one set in front of the ALU and the other set in front of the ID/EX register. The ID/EX muxes are only used for forwarding unsigned integers encountered during hazards with the LUI instruction. To get rid of this problem I changed the way my LUI instruction was handled. Instead of doing it inside of either the decoder or through its own logic, I passed the value straight into the ALU. If the LUI value goes into the ALU like a normal logical instruction any hazards encountered will be dealt with using the same hardware and logic used for the rest of the hazards, with the set of muxes in front of the ALU taking care of forwarding.

My initial hazard table was incomplete, as I neglected to add certain instruction combinations that would produce hazards. I had to add all of the *lui* instructions as well as the branching and jump instructions. I also made various mistakes on the correct forwarding logic for certain hazards. Below is the hazard table I used to design my forwarding and stalling units. This was provided to the class.

	R	AI	L	S	JR	LUI	B	JAL	
R	1a: $exmem.rd=idex.rs(1)$ $=idex.rt(2)$ 2a: $memwb.rd=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rd=ifd.rs(5)$ $=ifd.rt(6)$	1a: $exmem.rd=idex.rs(1)$ 2a: $memwb.rd=idex.rs(3)$ 3a: $memwb.rd=ifd.rs(5)$	1a: $exmem.rd=idex.rs(1)$ 2a: $memwb.rd=idex.rs(3)$ 3a: $memwb.rd=ifd.rs(5)$	1a: $exmem.rd=idex.rs(1)$ $=idex.rt(2)$ 2a: $memwb.rd=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rd=ifd.rs(5)$ $=ifd.rt(6)$	1a: $idex.rd=ifd.rs(stall)$ 2a: $exmem.rd=ifd.rs(8)$ 3a: $memwb.rd=ifd.rs(5)$		1a: $idex.rd=ifd.rs(stall)$ $=ifd.rt(stall)$ 2a: $exmem.rd=ifd.rs(8)$ $=ifd.rt(10)$ 3a: $memwb.rd=ifd.rs(5)$ $=ifd.rt(6)$		
AI	1a: $exmem.rt=idex.rs(1)$ $=idex.rt(2)$ 2a: $memwb.rt=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$	1a: $exmem.rt=idex.rs(1)$ 2a: $memwb.rt=idex.rs(3)$ 3a: $memwb.rt=ifd.rs(5)$	1a: $exmem.rt=idex.rs(1)$ 2a: $memwb.rt=idex.rs(3)$ 3a: $memwb.rt=ifd.rs(5)$	1a: $exmem.rt=idex.rs(1)$ $=idex.rt(2)$ 2a: $memwb.rt=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$	1a: $idex.rt=ifd.rs(stall)$ 2a: $exmem.rt=ifd.rs(8)$ 3a: $memwb.rt=ifd.rs(5)$		1a: $idex.rt=ifd.rs(stall)$ $=ifd.rt(stall)$ 2a: $exmem.rt=ifd.rs(8)$ $=ifd.rt(10)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$		
L	1a: $ifd.rt=pqe.rs(stall)$ $=pqe.rt(stall)$ 2a: $memwb.rt=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$	1a: $ifd.rt=pqe.rs(stall)$ 2a: $memwb.rt=idex.rs(3)$ 3a: $memwb.rt=ifd.rs(5)$	1a: $ifd.rt=pqe.rs(stall)$ 2a: $memwb.rt=idex.rs(3)$ 3a: $memwb.rt=ifd.rs(5)$	1a: $ifd.rt=pqe.rs(stall)$ $=pqe.rt(stall)$ 2a: $memwb.rt=idex.rs(3)$ $=idex.rt(4)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$	1a: $ifd.rt=pqe.rs(stall)$ 2a: $idex.rt=pqe.rs(stall)$ 3a: $memwb.rt=ifd.rs(5)$	1a: $ifd.rt=pqe.rs(stall)$ 2a: $idex.rt=pqe.rs(stall)$ 3a: $memwb.rt=ifd.rs(5)$		1a: $ifd.rt=pqe.rs(stall)$ $=pqe.rt(stall)$ 2a: $idex.rt=pqe.rs(stall)$ $=pqe.rt(stall)$ 3a: $memwb.rt=ifd.rs(5)$ $=ifd.rt(6)$	
S									
JR									
LUI	1a: $exmem.rt=idex.rs(11)$ $=idex.rt(12)$ 2a: $memwb.rt=idex.rs(13)$ $=idex.rt(14)$ 3a: $memwb.rt=ifd.rs(15)$ $=ifd.rt(16)$	1a: $exmem.rt=idex.rs(11)$ 2a: $memwb.rt=idex.rs(13)$ 3a: $memwb.rt=ifd.rs(15)$	1a: $exmem.rt=idex.rs(11)$ 2a: $memwb.rt=idex.rs(13)$ 3a: $memwb.rt=ifd.rs(15)$	1a: $exmem.rt=idex.rs(11)$ $=idex.rt(12)$ 2a: $memwb.rt=idex.rs(13)$ $=idex.rt(14)$ 3a: $memwb.rt=ifd.rs(15)$ $=ifd.rt(16)$	1a: $idex.rt=ifd.rs(17)$ 2a: $exmem.rt=ifd.rs(18)$ 3a: $memwb.rt=ifd.rs(15)$		1a: $idex.rt=ifd.rs(17)$ $=ifd.rt(19)$ 2a: $exmem.rt=ifd.rs(18)$ $=ifd.rt(20)$ 3a: $memwb.rt=ifd.rs(15)$ $=ifd.rt(16)$		
B									
JAL	1a: $exmem.31=idex.rs(21)$ $=idex.rt(22)$ 2a: $memwb.31=idex.rs(23)$ $=idex.rt(24)$ 3a: $memwb.31=ifd.rs(25)$ $=ifd.rt(26)$	1a: $exmem.31=idex.rs(21)$ 2a: $memwb.31=idex.rs(23)$ 3a: $memwb.31=ifd.rs(25)$	1a: $exmem.31=idex.rs(21)$ 2a: $memwb.31=idex.rs(23)$ 3a: $memwb.31=ifd.rs(25)$	1a: $exmem.31=idex.rs(21)$ $=idex.rt(22)$ 2a: $memwb.31=idex.rs(23)$ $=idex.rt(24)$ 3a: $memwb.31=ifd.rs(25)$ $=ifd.rt(26)$	1a: $idex.31=ifd.rs(27)$ 2a: $exmem.31=ifd.rs(28)$ 3a: $memwb.31=ifd.rs(25)$	1a: $idex.31=ifd.rs(27)$ $=ifd.rt(29)$ 2a: $exmem.31=ifd.rs(28)$ $=ifd.rt(30)$ 3a: $memwb.31=ifd.rs(25)$ $=ifd.rt(26)$			

Figure 1 - Hazard table

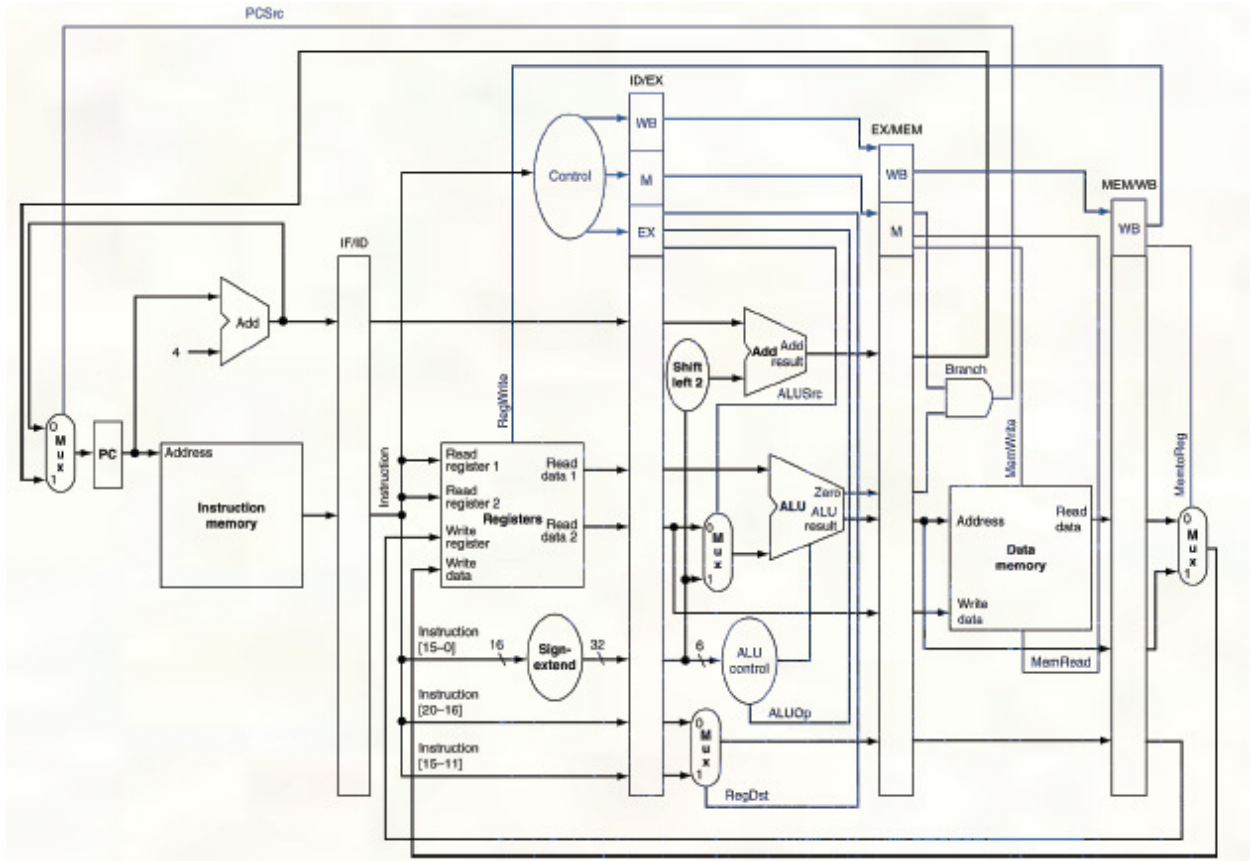


Figure 2 - Simplified view of the pipelined datapath and stages

Book Questions

Using the blue book.

4.12.1

a) Pipelined = 500ps, non-pipelined = $(500+100+350+400+300) = 1650\text{ps}$

b) Pipelined = 200ps, non-pipelined = $(200+150+120+190+140) = 800\text{ps}$

4.12.2

a) Pipelined = 1650ps (with forwarding), non-pipelined = 1650ps

b) Pipelined = 800ps (with forwarding), non-pipelined = 800ps

4.12.3

Split the highest latency stage

a) MEM, new time = 400ps

b) IF, new time = 190ps

4.20.1

a	<ul style="list-style-type: none"> lw followed by add (2 apart) lw followed by sw (3 apart) add followed by sw (1 apart) add followed by add (1 apart) add followed by sw (2 apart)
b	<ul style="list-style-type: none"> add followed by sw (1 apart) add followed by lw (2 apart) sw followed by lw (1 apart) sw followed by add (2 apart)

4.20.2

a	<ul style="list-style-type: none"> • With forwarding <ul style="list-style-type: none"> ○ none • Without forwarding <ul style="list-style-type: none"> ○ lw followed by add (2) ○ add followed by add (1) ○ add followed by sw (2) ○ add followed by sw (1)
b	<ul style="list-style-type: none"> • With <ul style="list-style-type: none"> ○ lw followed by add (1) • Without forwarding <ul style="list-style-type: none"> ○ add followed by sw (1) ○ lw followed by add (1)

4.24.1

- a) Always taken = 75%, never taken = 25%
- b) Always taken = 60%, never taken = 40%

4.24.2

- a) 0%
- b) 25%

2.24.3

- a) 75%
- b) 60%

Section 3: Instruction

JR

Taken from lab4demo.mif file.

Instruction:

```
lui    7,0x0000  
ori    7,7,0x308  
jr     7
```

This will store 0x00000308 into register 7 and then jump to this address. Forwarding is needed because register 7 does not have the entire 0x00000308 value when JR needs it, so lui is forwarded for the ori instruction, and then that is forwarded for the JR instruction.

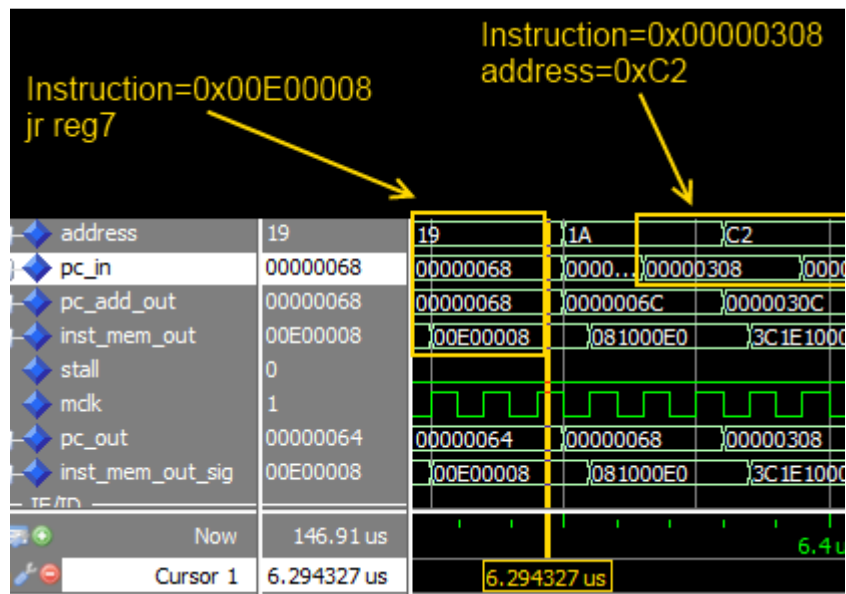


Figure 3 - JR going to contents of register 7

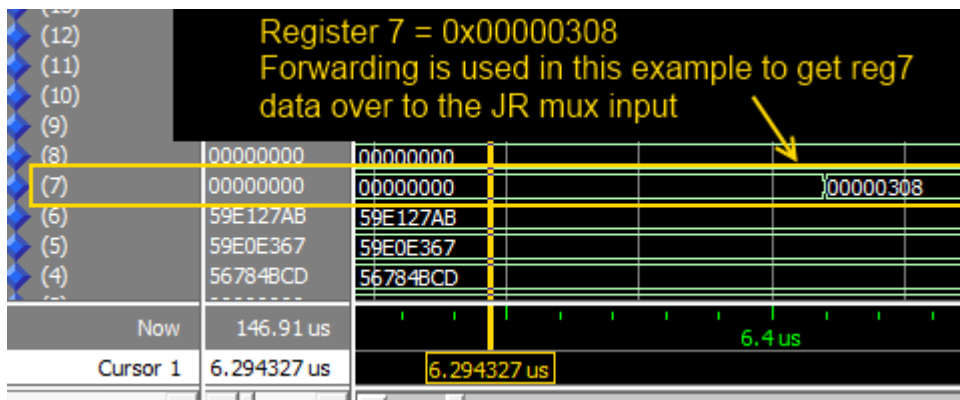


Figure 4 - JR, showing contents of register 7

Logical Functions

Test Code MIF	
lui	\$a0, \$zero, 0x2568
lui	\$a1, \$zero, 0x3a97
ori	\$a0, \$a0, 0xabef
ori	\$a1, \$a1, 0x5be0
and	\$a2,\$a0,\$a1
andi	\$a3,\$a0,0x4510
or	\$s0,\$a0,\$a1
nor	\$s1,\$a0,\$a1
sll	\$s2,\$a0,0x4
srl	\$s3,\$a0,0x2

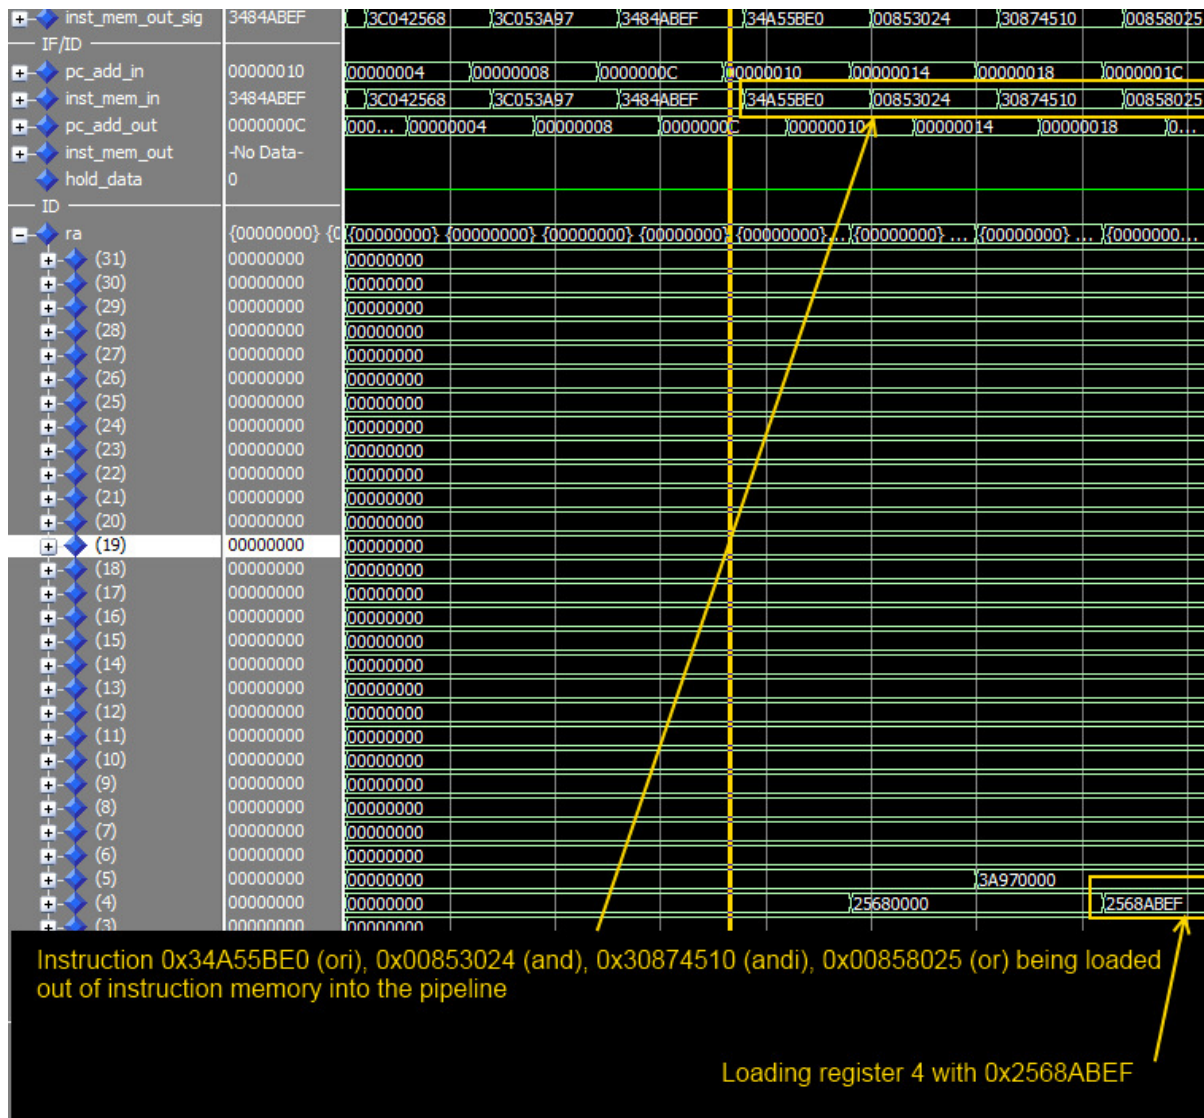


Figure 5 - Logical function test

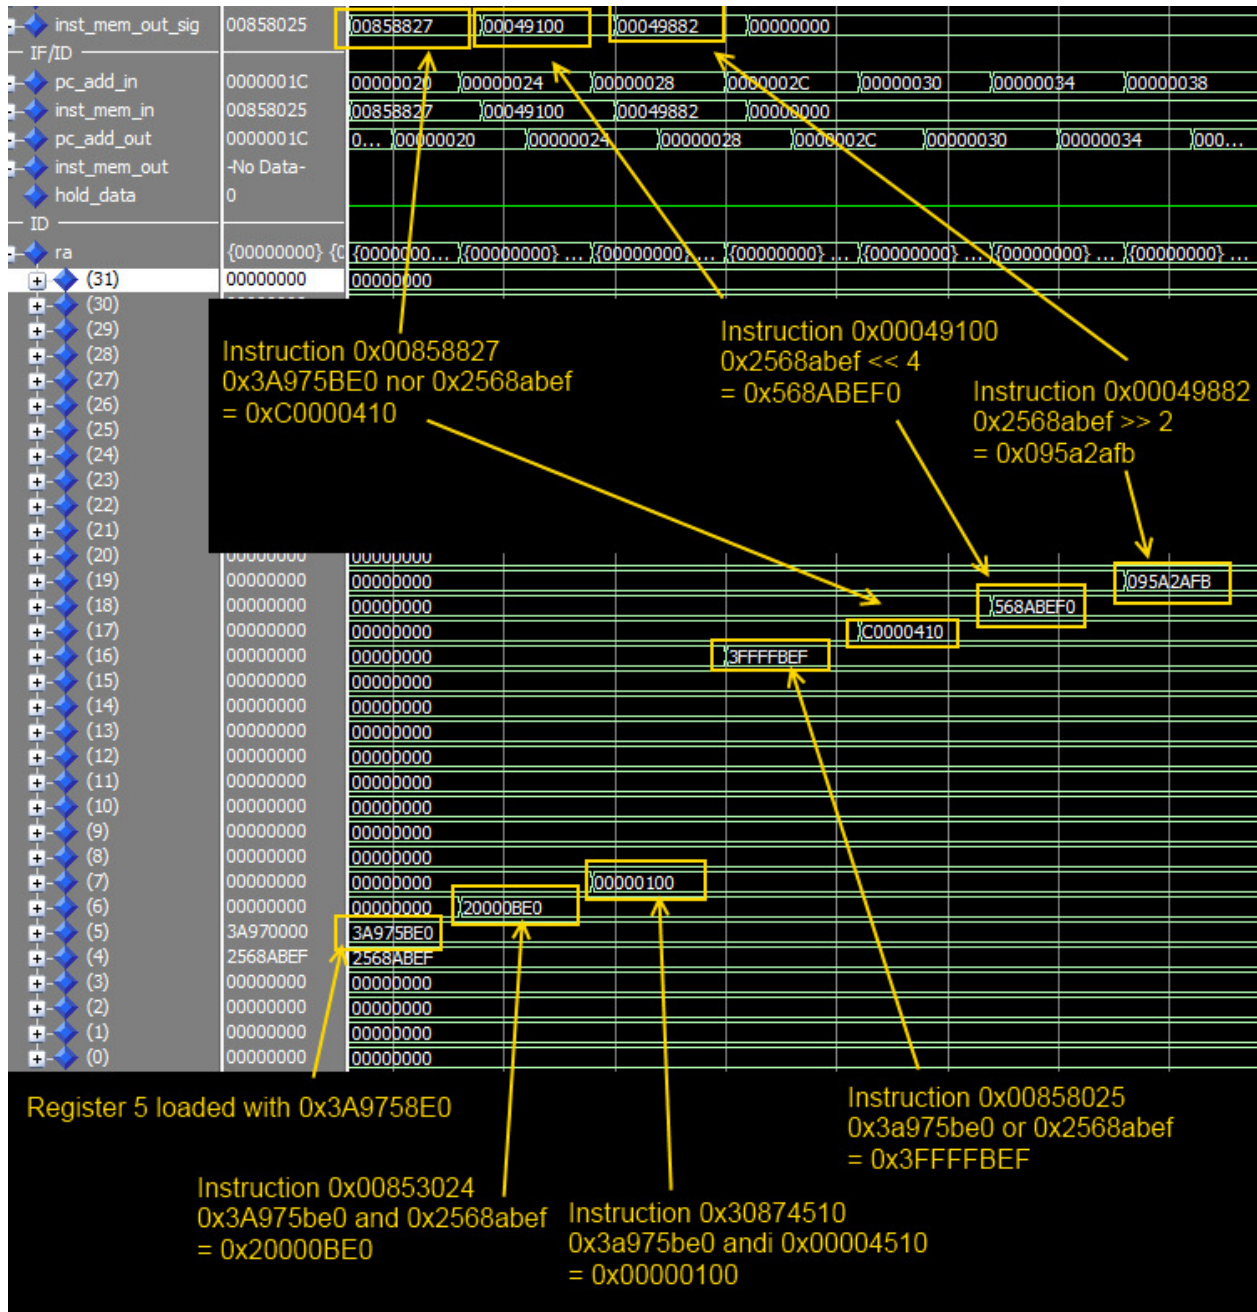


Figure 6 - Logical function test

BEQ/BNE

Test Branch MIF	
ori \$s0,\$0,0x1000	branch1:
ori \$s1,\$0,0x2000	bne \$s0,\$s2,end
ori \$s2,\$0,0x1000	bne \$s0,\$a1,end
beq \$s0,\$s1,branch1	lui \$s5,\$zero,0xFFFF
beq \$s0,\$s2,branch1	end:
lui \$s4,\$zero,0xFFFF	lui \$s6,\$zero,0xFFFF

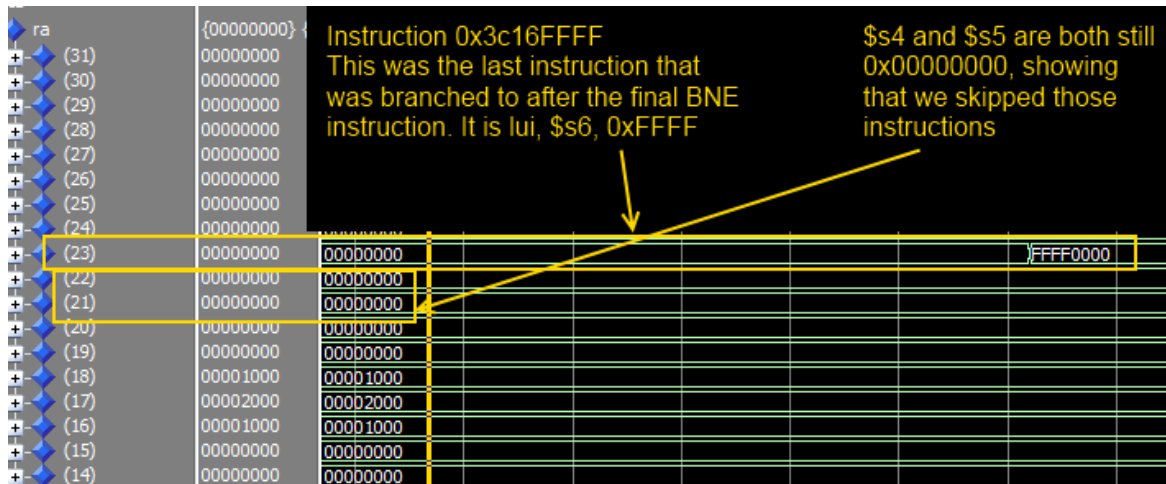


Figure 8 - Branch testing outcome

SLT, SLTU, SLTI, SLTIU

```

SLT/SLTU/SLTI/SLTIU Test MIF
lui    $a0,0xFF11
lui    $s0,0xFFFF
ori    $s0,$s0,0xFFFF
ori    $s1,$s1,0x1234
ori    $s2,$s2,0x2345
slt    $t0,$s1,$s2
slt    $t1,$s2,$s1
sltu   $t3,$s0,$s1
sltu   $t4,$s1,$a0
slti   $t5,$s1,0x4000
slti   $t6,$s2,0x1900
sltiu  $t7,$a0,-6
sltiu  $t8,$a0,0x4161

```

The difference between SLT and SLTU as well as SLTI and SLTIU is how the code perceives the data. Because my main example images didn't show the differences well enough I created another short MIF file to test the signed and unsigned version against two identical numbers. This can be seen and explained in the image below. The same concept applies to SLTI vs SLTIU as well.

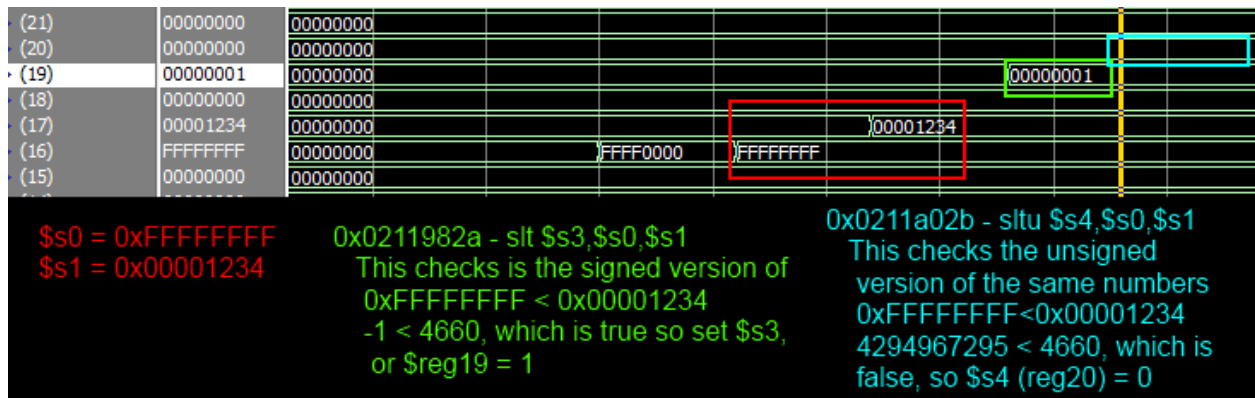


Figure 9 - Signed vs Unsigned

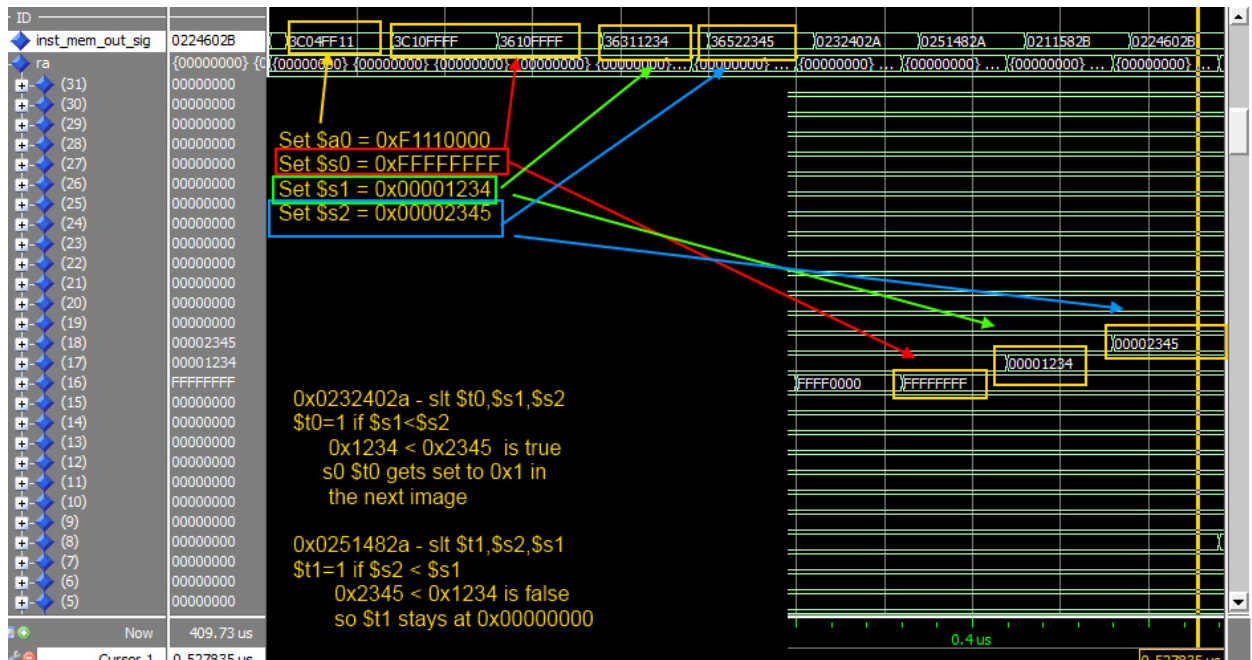


Figure 10 - SLT testing

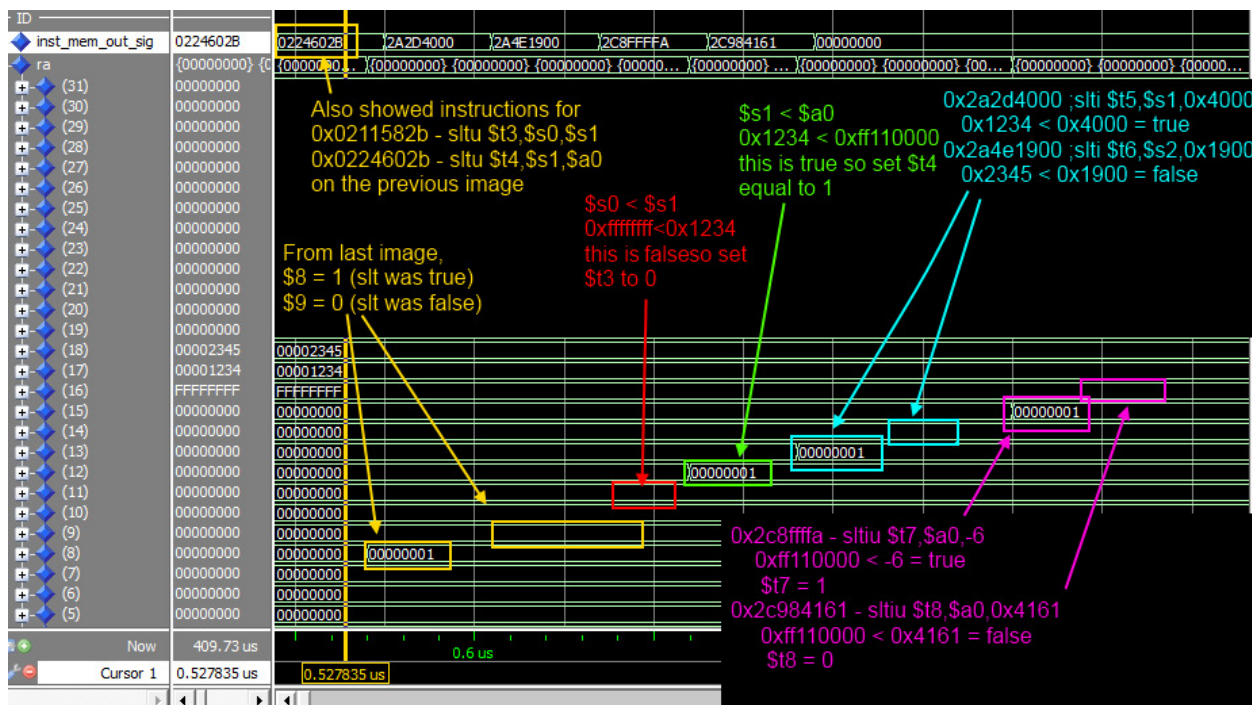


Figure 11 - SLT testing

SUB, SUBU, ADD, ADDU

The main difference between the ADD vs ADDU and SUB vs SUBU is that both of signed versions produce overflow while the unsigned version do not.

```

Test MIF code
lui    $s0,0x8193
addi   $s1, $s1,0x2468
addi   $s2,$s2,0x1234
add    $t0,$s1,$s2
addu   $t1,$s0,$s1
addu   $t2,$s1,$s2
sub    $t3,$s1,$s0
sub    $t4,$s1,$s2
subu   $t5,$s1,$s0
subu   $t6,$s1,$s2

```

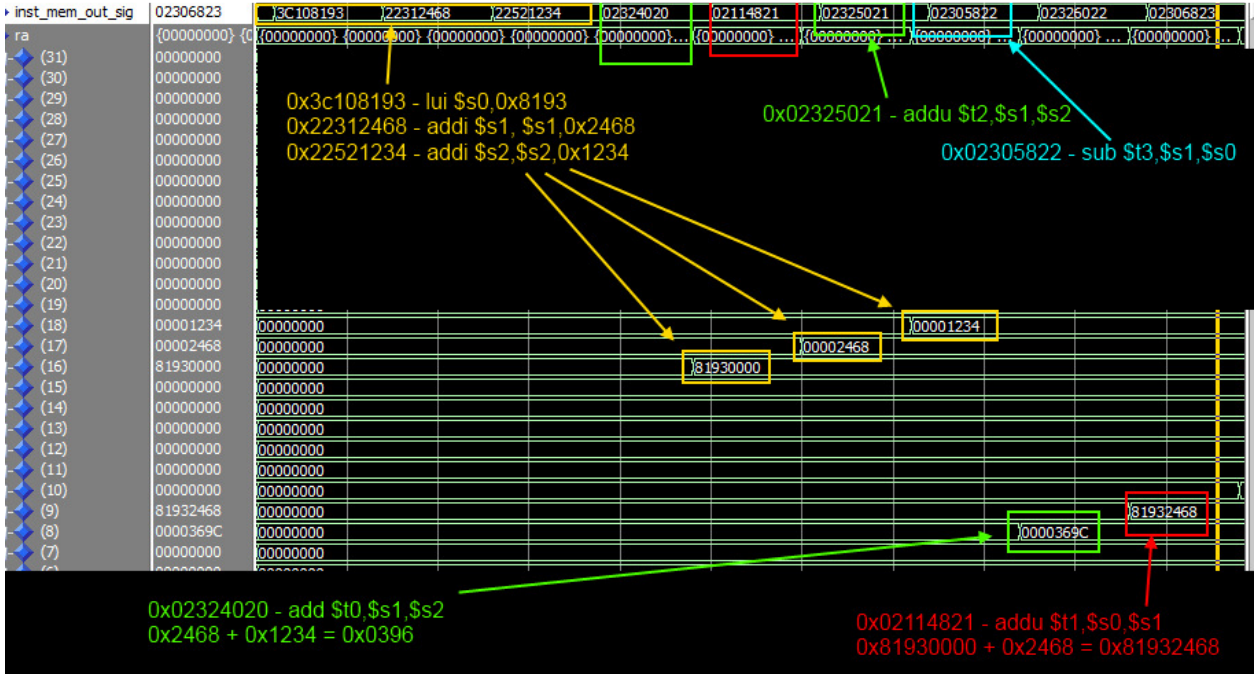


Figure 12 - SUB/SUBU/ADD/ADDU

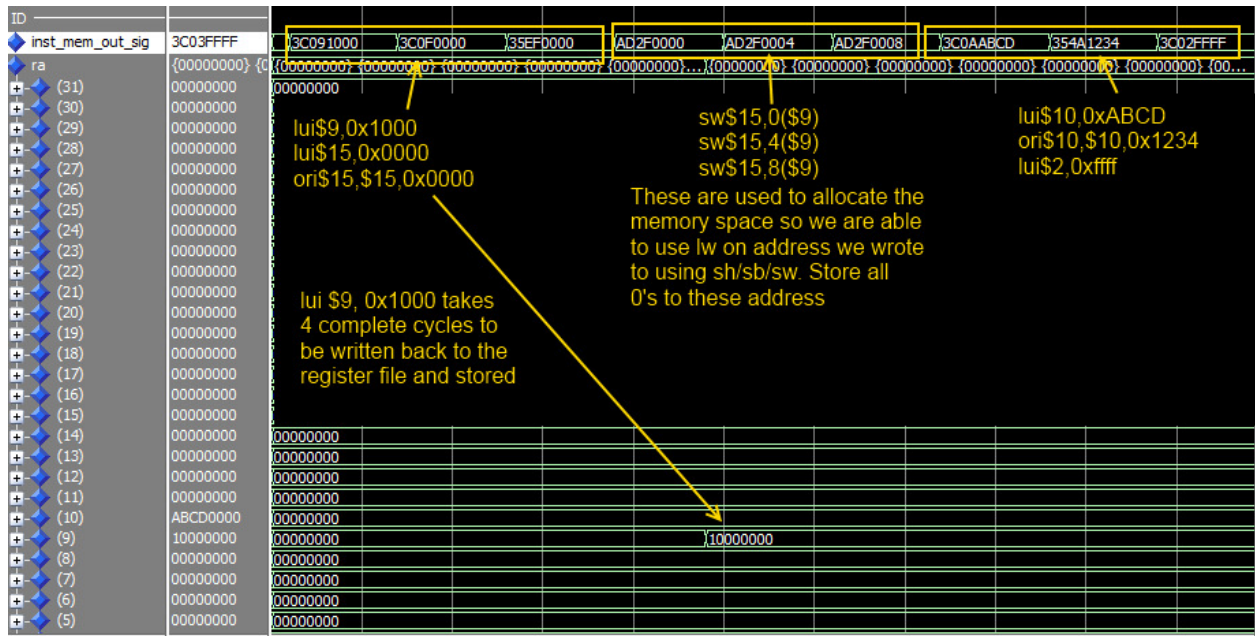
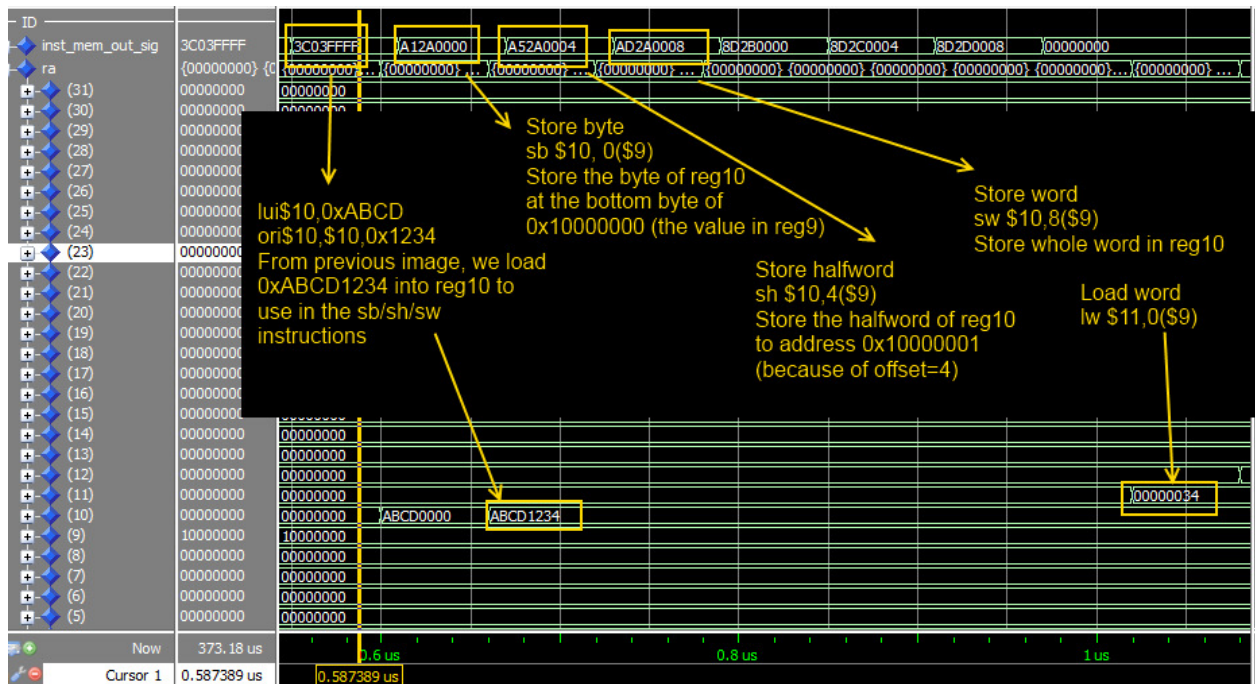


Figure 14 - Setting up the registers for the store test



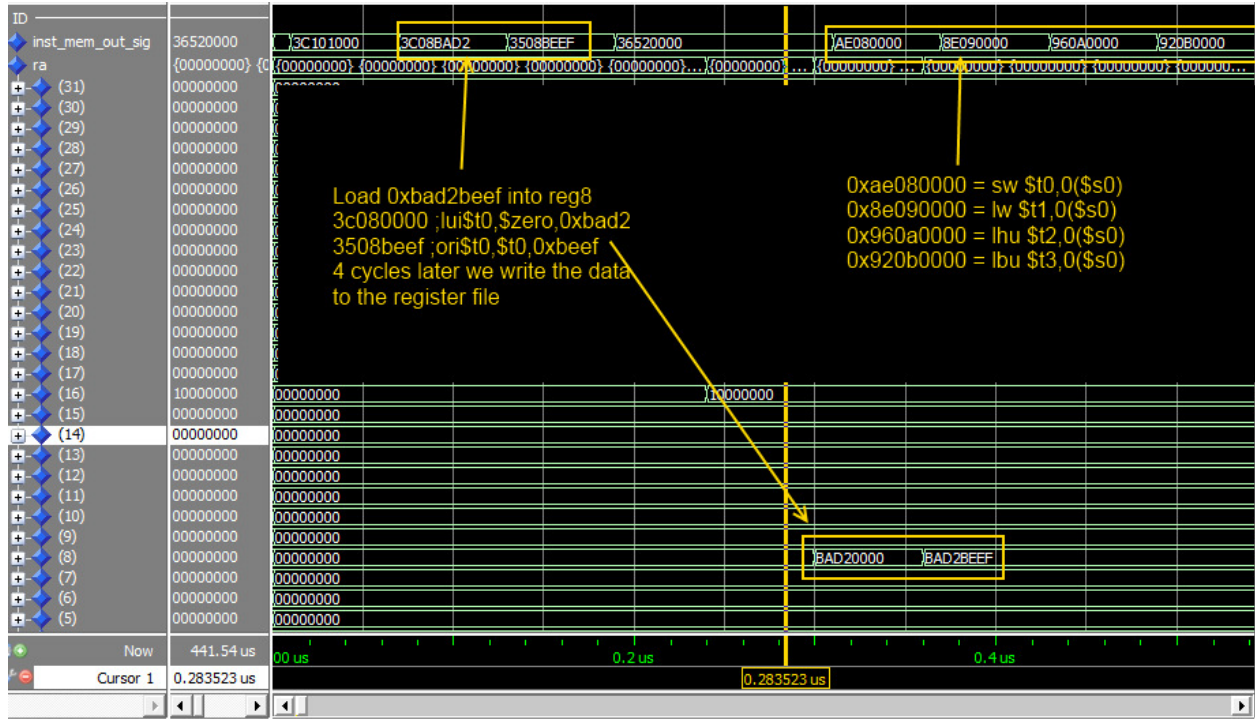


Figure 17 - Setting up registers to test LW, LHU, and LBU. Also shows the instructions in the IF stage

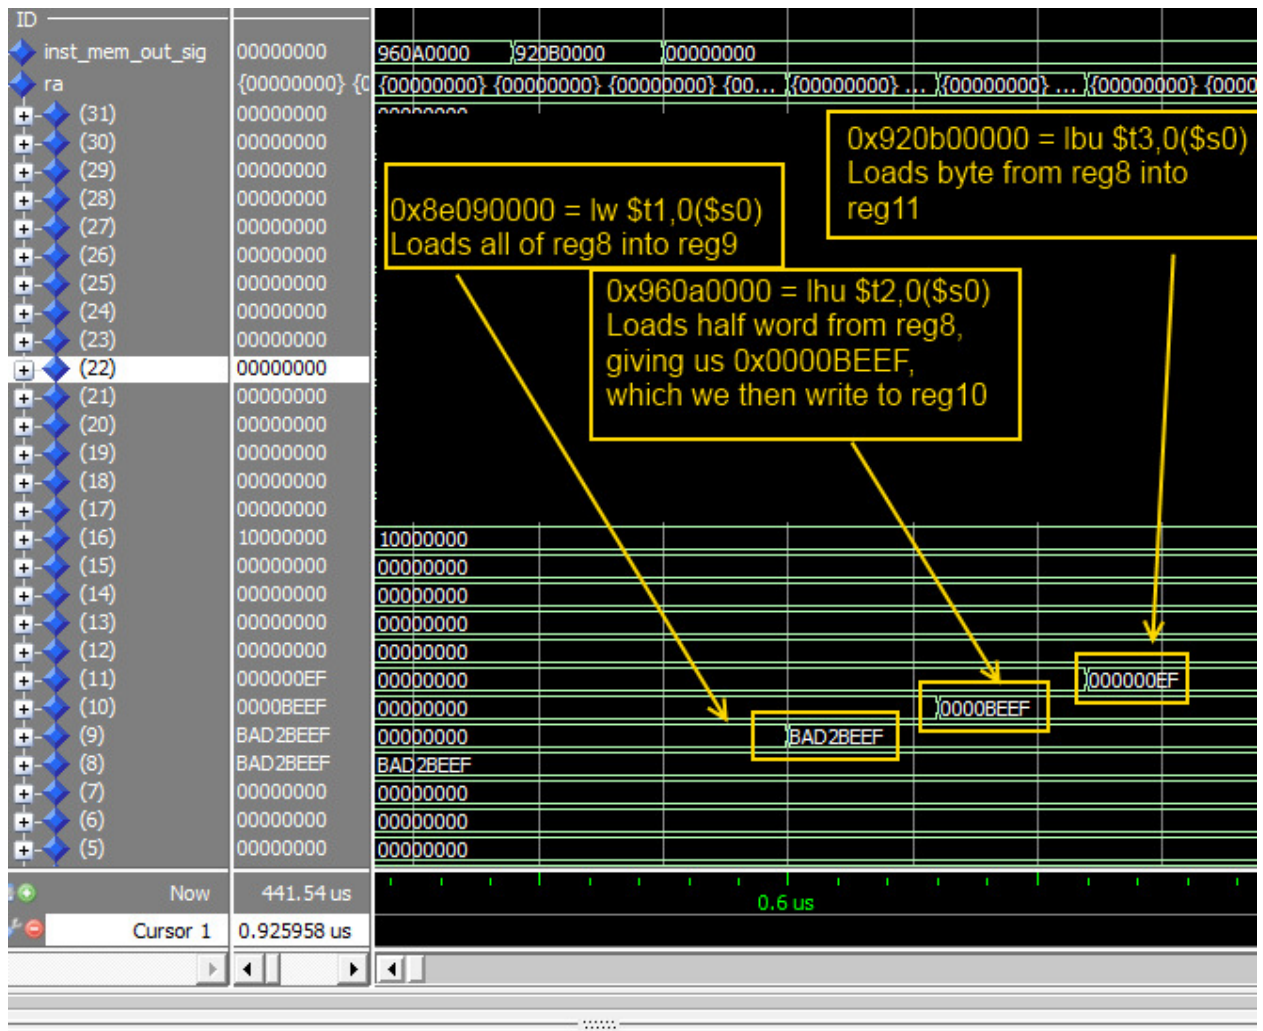
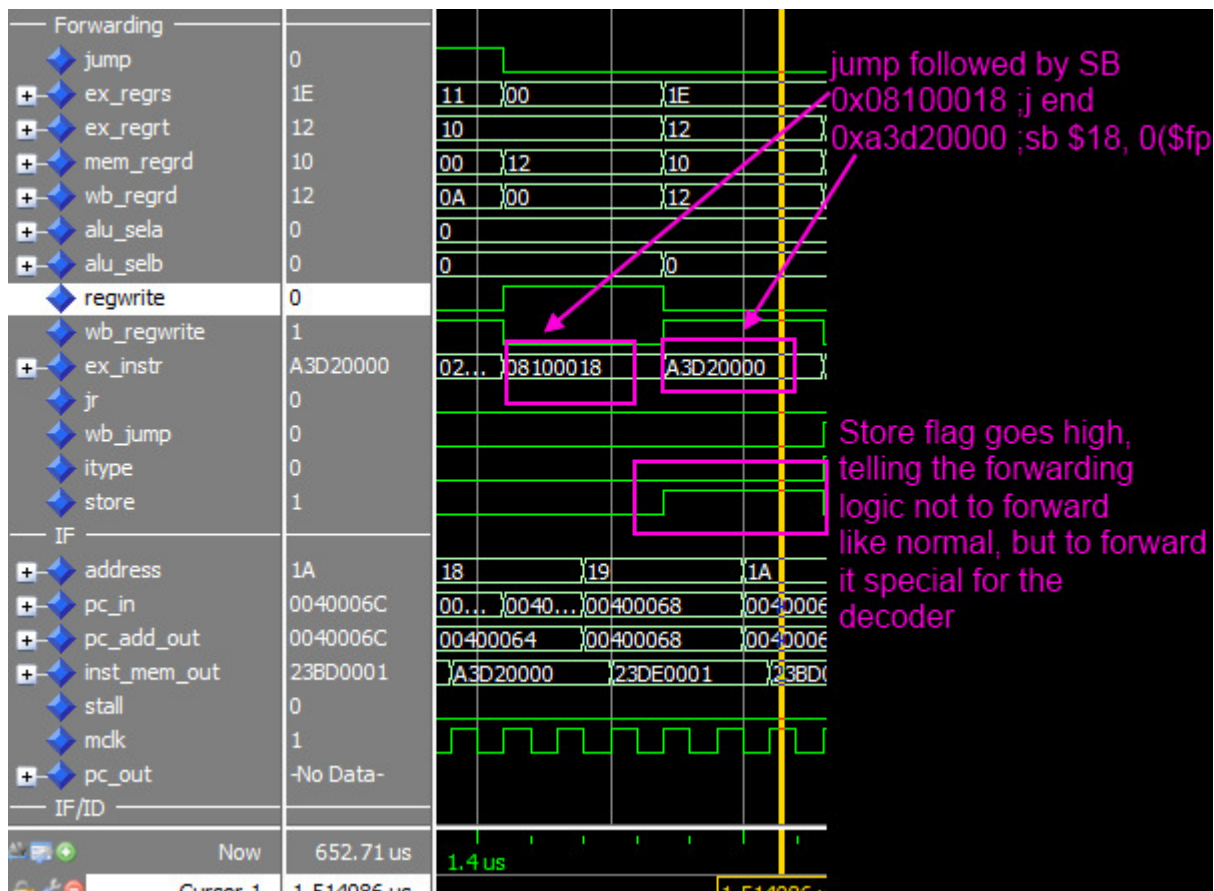


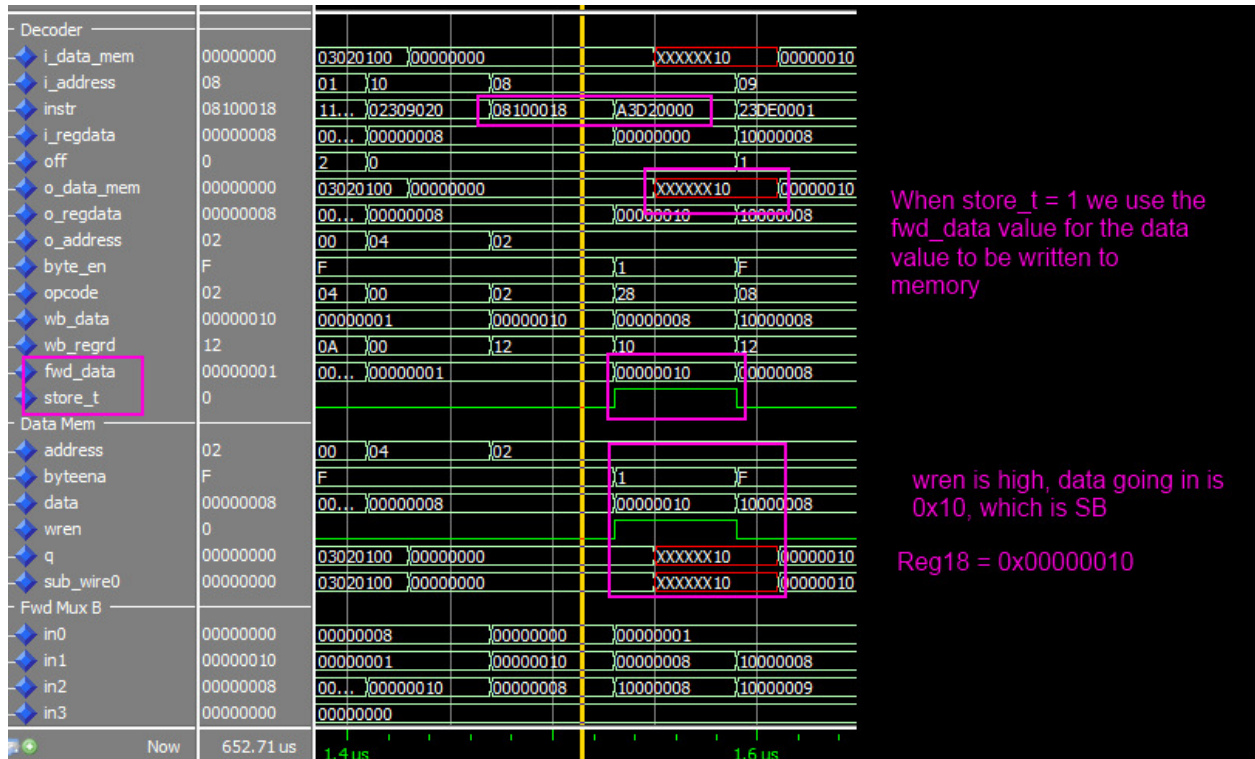
Figure 18 - Register values from the result of LW, LHU, and LBU

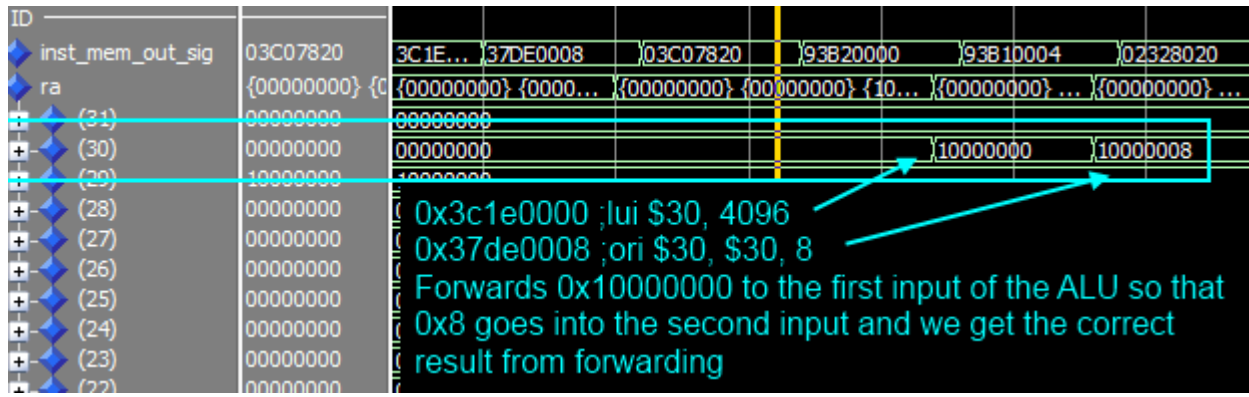
Section 4: Hazard detection and forwarding

I got this working 2 hours before the extended deadline, sorry for the quick images and short descriptions, I didn't plan to be working on it this long.

I had issues forwarding with my decoder so even though the forwarding worked fine for normal instructions, it wasn't working 100% with load and store instructions. The way I did forwarding for those was inside the actual decoder and partially inside the forward unit as well. In the forward unit I watched to see if the last instruction had a hazard (so if $ex.r\text{t}/ex.r\text{s} = mem.r\text{d}$). If so I sent the data from that register that needed to be forwarded along with a 1 bit flag signal through the pipeline, into MEM stage, and into the decoder. If the flag went true I took that forwarded data as the input, otherwise I took the normal data from the pipeline. I also routed the $wb.r\text{d}$ as well as the $wb.data$ into the decoder. I watched this to see if $mem.r\text{d} = wb.r\text{d}$, and if it did I forwarded, or used the $wb.data$ I had brought over. This was not a very fun or clean way to do this, but the traditional ways were giving me a lot of problems. Everytime I would fix one issue another would come up and I was already behind from other problems with the pipelining and previous checkpoints.







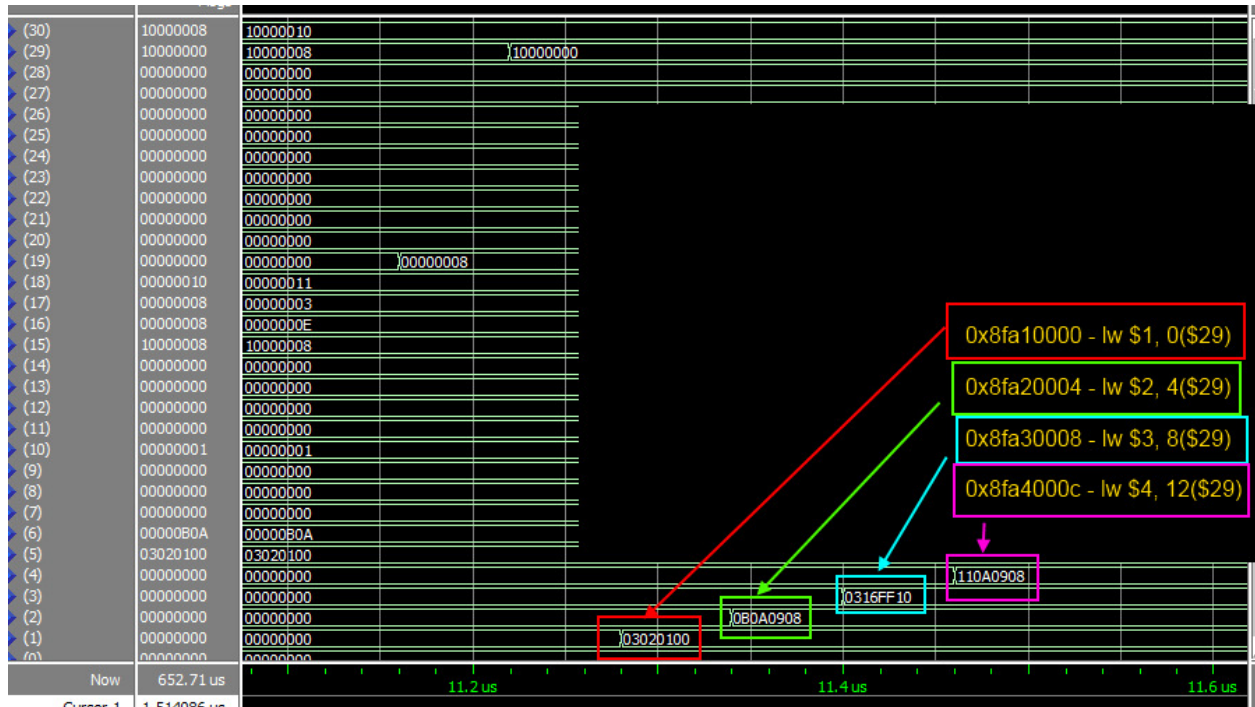


Figure 19 - Values of final registers

The program ends its execution by going into an infinite loop at the very end. A jump instruction keeps the loop going forever. The hazard mif file, when working correctly, outputs the following values

- \$1 = 0x03020100
- \$2 = 0x0b0a0908
- \$3 = 0x0316ff10
- \$4 = 0x110a0908

Full Register output

