Computer System Architecture
6.823 Quiz #1
March 10th, 2017
Professors Daniel Sanchez and Joel Emer

Name: _______**Solutions**__________

This is a closed book, closed notes exam.
80 Minutes
18 Pages (+2 Scratch)

Notes:
- Not all questions are of equal difficulty, so look over the entire exam and budget your time carefully.
- Please carefully state any assumptions you make.
- Show your work to receive full credit.
- Please write your name on every page in the quiz.
- You must not discuss a quiz's contents with other students who have not yet taken the quiz.
- Pages 19 and 20 are scratch pages. Use them if you need more space to answer one of the questions, or for rough work.

Part A	________	25 Points
Part B	________	35 Points
Part C	________	40 Points
TOTAL	________	**100 Points**

Part A: Self-modifying Code (25 points)

In this question, you will implement linked-list operations using self-modifying code on an EDSACjr machine. The memory layout is shown in the figure on the right. You have access to the named memory locations as indicated. Linked-list nodes consist of two words: the first is an integer `value`, the second is an address pointing to the `next` node. `_HEAD` contains the address of the first node of the list (or `_INVALID` if it is empty). The `next` field of the last node is `_INVALID`. All valid addresses are positive. You may create new local and global labels as explained in the EDSACjr handout.

Table A-1 shows the EDSACjr instruction set.

Opcode	**Description**	**Bit Representation**
ADD *n*	Accum ← Accum + M[*n*]	00001 *n*
SUB *n*	Accum ← Accum - M[*n*]	10000 *n*
STORE *n*	M[n] ← Accum	00010 *n*
CLEAR	Accum ← 0	00011 00000000000
OR *n*	Accum ← Accum | M[*n*]	00000 *n*
AND *n*	Accum ← Accum & M[*n*]	00100 *n*
SHIFTR *n*	Accum ← Accum shiftr *n*	00101 *n*
SHIFTL *n*	Accum ← Accum shiftl *n*	00110 *n*
BGE *n*	If Accum ≥ 0 then PC ← *n*	00111 *n*
BLT *n*	If Accum < 0 then PC ← *n*	01000 *n*
END	Halt machine	01010 00000000000

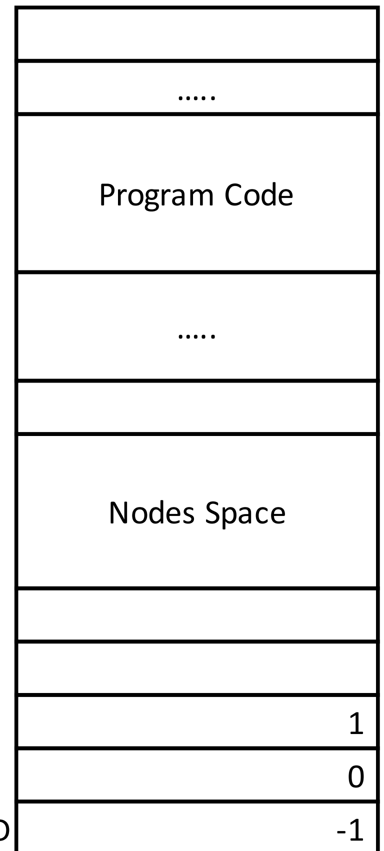

You may also use the following macros if required.

Macro	**Description**
STOREADR *n*	Replace the address field of location *n* with the contents of the accumulator
LOADADR *n*	Load the address field of location *n* into the accumulator

Question 1 (10 points)

Write a macro for **LISTPUSH**, which pushes the node pointed to by the accumulator to the head of the list. **LISTPUSH** takes one argument, the memory address of the new node, which is available in the accumulator. As shown in the figure below, **LISTPUSH** stores the current `_HEAD` pointer in the new node's `next` field, and updates the `_HEAD` pointer to point to the new node. Implement the macro using the EDSACjr instruction set and macros provided above. Do not refer to "`value`" or "`next`"; they are for illustration only. You need not worry about memory allocation; the new node's address is provided in the accumulator.

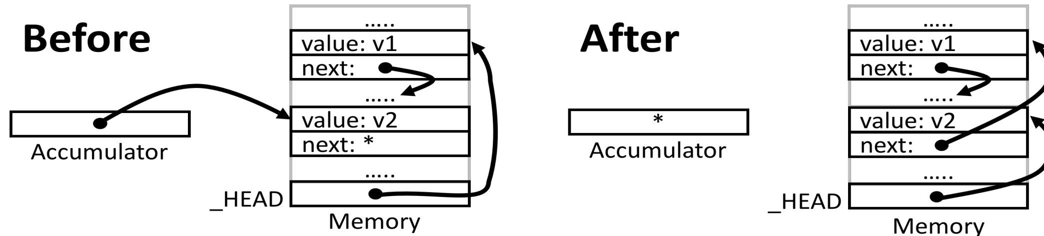


```
.macro LISTPUSH
      STORE _TMP        ;; store accumulator (address of the new node)
      ADD _ONE          ;; accum <- address of the new node's next field
      STOREADR _STN     ;; address field of location _STN has the address
                        ;; of the new node's next field
      CLEAR
      ADD _HEAD         ;; accum <- M[_HEAD], current head pointer
_STN: STORE 0           ;; 0 will be replaced with the node's next field
                        ;; address. M[_TMP + 1] <- accum
      CLEAR
      ADD _TMP          ;; retrieve address of new node in accumulator
      STORE _HEAD       ;; M[_HEAD] <- accum; Update the head pointer
                        ;; to the new node




.end
```

Question 2 (10 points)

Write a macro for **LISTPOP**, which removes the node at the head of the list and stores its address in the accumulator, or stores `_INVALID` (-1) in the accumulator if the list is empty. Implement the macro using the EDSACjr instruction set and macros provided above.

```
.macro LISTPOP
      CLEAR             ;; accumulator is not an input
      ADD _HEAD         ;; accum <- address of head node
      BLT _DONE         ;; if _HEAD < 0 (-1, ie *_INVALID), then return
      STORE _TMP        ;; save old value of head
      ADD _ONE          ;; accum <- address of head node's next field
      STOREADR _ADDN    ;; replace address field of _ADDN
                        ;; with address of head node's next field
      CLEAR
_ADDN: ADD 0            ;; 0 will be replaced with the address of head
                        ;; node's next field. accum <- addr of 2nd node
      STORE _HEAD       ;; update head with the list's second node
      CLEAR
      ADD _TMP          ;; accum <- former head node pointer
_DONE:















.end
```

Question 3 (5 points)

Assume there exists a macro called **FREE** that takes an address as input in the accumulator and deallocates it (just like `free(void* ptr)` in C). Write a macro for **LISTCLEAR**, which uses the **FREE** macro and your **LISTPOP** macro to remove and deallocate all nodes in the list. Assume all valid node addresses are positive, or else a pointer is `_INVALID` (-1). Implement the macro using the EDSACjr instruction set and macros provided above.

```
.macro LISTCLEAR
_LOOP:     LISTPOP          ; accum <- address of removed node
           BLT _DONE        ; exit if an _INVALID node pointer is found
           FREE             ; de-allocates the removed node
           CLEAR
           BGE _LOOP
_DONE:

.end
```

Part B: Virtual Memory

Ben Bitdiddle purchases a new processor to run his 6.823 labs. The processor manual indicates that the machine is byte-addressed with 20-bit virtual addresses and 20-bit physical addresses. The following figure summarizes the 2-level page table structure and shows the breakdown of a virtual address in this system. The physical address of the base of the Level 1 page table (0x02000) is stored in the L1 Table Base Address register. The L1 and L2 page tables are located in physical memory. The size of both L1 and L2 page table entries is 4 bytes. Each entry of the L1 page table contains the physical address of the base of each Level 2 page table (a PTP), and each of the L2 page table entries holds the PTE of the data page.

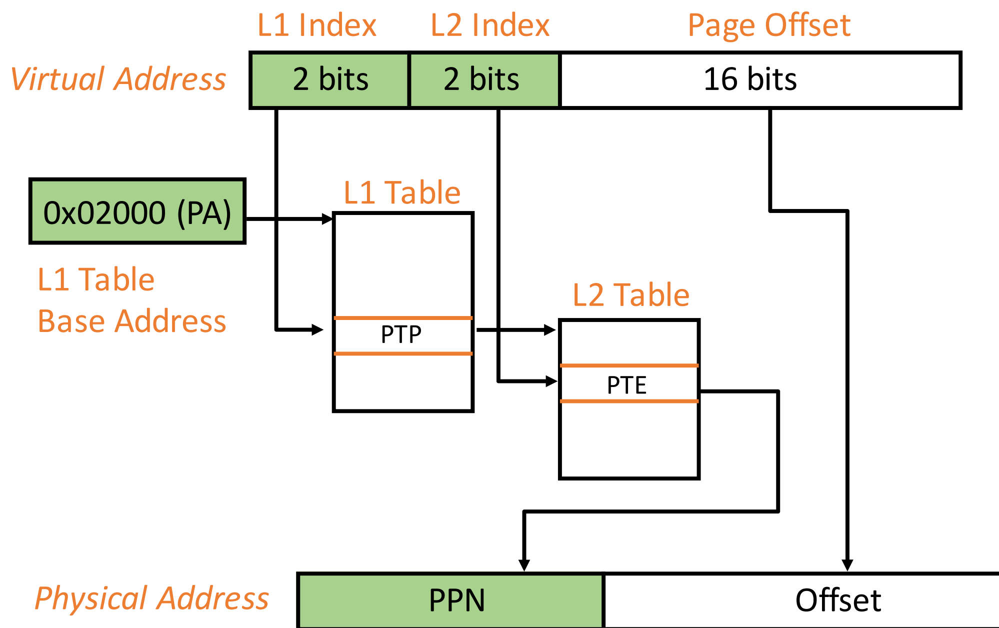

A PTE in L2 page tables can be broken into the following fields. (Don't worry about status bits).

31 — 20	19 — 16	15 — 0
0	Physical Page Number (PPN)	Status Bits

A PTP in the L1 page table appears as follows.

31 — 20	19 — 0
0	Physical address of a L2 page table

Question 1 (10 points)

Assuming the initial memory state is as shown to the right, what is the physical page number (PPN) of virtual address (VA) 0xB29A0? What is the physical address (PA)? Show and explain your work for full credit. **For your convenience, we separate the page number from the offset with a colon ":".**

Address (PA)	
0x0:2000	0x0:2048
0x0:2004	0x0:2010
0x0:2008	0x0:2038
0x0:200C	0x0:2028
0x0:2010	0x1:0084
0x0:2014	0x5:0DA8
0x0:2018	0x6:11A0
0x0:201C	0xB:9944
0x0:2020	0xC:7FFF
0x0:2024	0x4:B000
0x0:2028	0x7:30B1
0x0:202C	0xD:2E5C
0x0:2030	0x3:A000
0x0:2034	0x6:010C
0x0:2038	0xA:74C0
0x0:203C	0x8:A524
0x0:2040	0x9:FFEE
0x0:2044	0x2:93A4
0x0:2048	0xA:74D0
0x0:204C	0x3:FD40

Snapshot of physical memory

<u>**Virtual Address**</u>

`0xB:29A0 = 0b 1011:0010100110100000`

The L1 base PA is 0x02000 (shown in the figure).
The L1 index of the given VA is 2 (0b10).
Since PTP entries are 32 bits (4 bytes), the L1 PTP appears at 0x02000 + 2x4 = 0x02008.
Reading that PA, the base PA of the L2 table is 0x02038.
The L2 index of the given VA is 3 (0b11).
Therefore the L2 PTE appears at 0x02038 + 3x4 = 0x02044.
Reading that PA, the L2 PTE is 0x293A4.
We ignore the 16 status bits and extract page number 2.

VPN 0xB => PPN _____**0x2**_________

The page offset from the VA doesn't change, so append the VA offset after PPN 2 to construct the physical address.

VA 0xB29A0 => PA ____**0x229A0**_________

Unable to run Pin in his own environment, Ben's friend, Alyssa P. Hacker, refers him to the Nested Paging handout to learn how to run his labs in a virtual machine (much to the TA's dismay!) However, Ben is frustrated by the worst-case performance. Let's find out why.

Question 2 (4 Points)

Ben starts his foray into virtualization by thinking about `gPA=>hPA` translation.

a) Assuming Ben's host physical memory has the same snapshot as in Question 1, what is the host physical address (`hPA`) of guest physical address (`gPA) 0xB29A4`? Explain.

As stated in the Nested Paging handout, a `gPA` appears to the host system as a host virtual address. Therefore, `gPA 0xB29A4` is an address 4 bytes higher than the (host) virtual address of Question 1, and its (host) `PPN` remains the same: 2. We append page offset `0x29A4` to that `PPN`.

`gPA 0xB29A4 => hPA ____0x229A4__________`

b) Assuming no TLB, how many accesses to host physical memory are required to access the data associated with a `gPA` (i.e., perform the `gPA=>hPA` translation and fetch the data)? Explain.

Three. No different than accessing the data associated with a host virtual address.

One indexed access into the host L1 table
One indexed access into the host L2 table, (i.e. two for gPA=>hPA translation)
One to fetch the data at the identified host physical address.

Question 3 (12 Points)

Given a guest virtual address (`gVA`), the first step of a nested page table walk is to load the relevant guest L1 page table PTP. This provides the base `gPA` of the guest L2 table. Ben is shocked at how much work is required!

a) Assume host physical memory is initialized as in Question 1 and as shown to the right, the Guest Table Base Address register holds `0xB29A0` (a `gPA`), and the Host Table Base Address register holds `0x02000` (a `hPA`). During a nested page table walk of guest virtual address (`gVA`) `0x61EAC`, what are the contents of the guest L1 page table PTP entry? **For your convenience, we separate the page number from the offset with a colon ":".**

Address (PA)	
0x0:2000	0x0:2048
0x0:2004	0x0:2010
0x0:2008	0x0:2038
0x0:200C	0x0:2028
0x0:2010	0x1:0084
...	...
0x2:2998	0xD:2E5C
0x2:299C	0x3:A000
0x2:29A0	0x6:010C
0x2:29A4	0xA:74CC
0x2:29A8	0x7:30B1
...	...
0x5:2994	0x6:11A0
0x5:2998	0xB:F149
0x5:299C	0xC:7BFF
0x5:29A0	0x4:B020
0x5:29A4	0xF:A120
...	...
0xA:299C	0x8:A624
0xA:29A0	0x9:FEED
0xA:29A4	0x2:93A4
0xA:29A8	0xA:7440
0xA:29AC	0x3:FD40

Snapshot of **host** physical memory

Guest Virtual Address (gVA)

`0x6:1EAC   =   0b 0110:0001111010101100`

The guest L1 base `gPA` is `0xB29A0`.
The L1 index of the given `gVA` is 1 (`0b01`).
PTP entries are 4 bytes, so the guest L1 PTP appears at
`gPA 0xB29A0 + 4 = 0xB29A4`.
From Question 2, this corresponds to `hPA 0x229A4`.
Reading that `hPA`, we find the contents of the Guest L1 Table PTP

Guest L1 Table PTP = ___**`0xA74CC`**_________

b) Assume no TLB. Starting from some `gVA`, how many accesses to host physical memory are required to determine the guest L1 PTP entry of a guest virtual address (`gVA`)? Explain.

Three.

We can determine the `gPA` of the guest L1 PTP entry with zero memory accesses:
add the `gVA` L1 index offset to the Guest Table Base Address (which appears in a register).

Given this `gPA` of the L1 PTP entry, we must load its data. From Q2b, 3 accesses to host physical memory are required to fetch the data of a `gPA`.

Question 4 (4 Points)

For the 2-level nested page table in the Nested Paging handout, assuming no TLB, how many accesses to host physical memory are required to perform a guest memory access? (i.e. given a gVA, find its corresponding hPA **and fetch the data**). Explain.

Nine.

Load the guest L1 PTP (from a gPA), whose data is the gPA of the guest L2 table's base.
Load the guest L2 PTE (from a gPA), whose data is the gPA of the data to be fetched.
Load the data (from a gPA).

Each load of these three gPAs requires 3 accesses to fetch the data.

Question 5 (5 Points)

For an M-level hierarchical guest page table and an N-level hierarchical host page table, assuming no TLB, how many accesses to host physical memory are required to perform a guest memory access? (i.e. given a gVA, find its corresponding hPA **and fetch the data**). Explain.

$(M+1)$ $(N+1)$

Generalizing Q2b, any gPA=>hPA translation requires walking the N levels of the host page table (which consists exclusively of hPAs, and therefore accesses to host physical memory). Given the data's hPA, one more access is required to fetch the data: $(N+1)$ accesses to fetch the data of a gPA.

The gVA=>gPA translation requires walking the M levels of the guest page table. The guest page table consists of gPAs. Each level of the walk requires accessing a gPA and fetching its data (to find the gPA address of the next level's base table), so the gVA=>gPA translation requires $M(N+1)$ accesses to host physical memory.

Given the final gPA of the gVA, one more gPA access is required to fetch the data, which requires $N+1$ accesses.

$M(N+1) + N+1$

Part C: ISA/Instruction Pipelining

In this problem we study a pipelined implementation of BigMIPS. The BigMIPS ISA is given as a handout.

Question 1: ISA Design (5 points)

The BigMIPS ISA uses consecutive 32-bit registers to form a 64-bit base address for Load Word (`LW`) and Store Word (`SW`) instructions. Could the `LW` and `SW` instruction formats incorporate an arbitrary pair of source registers to form the base address? If yes, are there any consequences? If no, why not? (One to two sentence response.)

"Yes" or "No" answers are acceptable with proper justification.

Yes: instead of being based on the Immediate-Type instruction format, `LW` and `SW` could use the Register-Type instruction format, which has three register fields, and interpret the `sa` and `funct` bit fields as an `offset`. This approach would reduce the size of the `offset` from 16 bits to 11.

No: loads and stores are of the Immediate-Type instruction format, and cannot support a third register without changing instruction decode logic. Additionally, this would reduce the size of the offset from 16 bits to 11 bits, which the quiz taker can deem unacceptable.

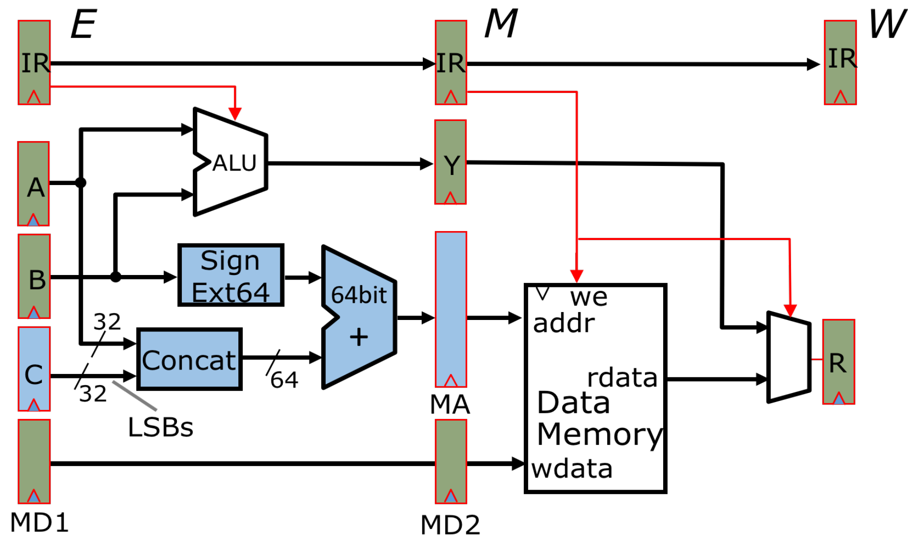

Figure C.1: Updates to the MIPS Execute and Memory stages to support the BigMIPS ISA.

Alyssa P. Hacker extends the classic 5-stage MIPS pipeline to implement the BigMIPS ISA. She first focuses on the Execute and Memory stages, shown in Figure C.1 (we'll see the Decode stage next). The figure highlights Alyssa's additions in blue. The figure shows how she builds a 64-bit effective address for `LW` and `SW` in the Execute stage. Alyssa assumes the following about Execute-stage registers:

- *A* holds (*rs*)
- *B* holds the 32-bit extended offset
- C holds (*rs+1*)
- *MD1* holds (*rt*), the data to store (**Do not worry about getting data for stores for now**)

Alyssa adds a 64-bit adder that sums the base address and 64-bit extended offset. Register *C* (i.e. (*rs+1*)) makes up the 32 least significant bits (LSBs) of the base address. At the memory stage, Alyssa adds a 64-bit Memory Address register (*MA*).

Question 2: Loads (7 points)

Alyssa updates the Decode stage with a mux before *rs2*, as shown below. She wants your help to fill registers *A*, *B*, and *C* correctly.

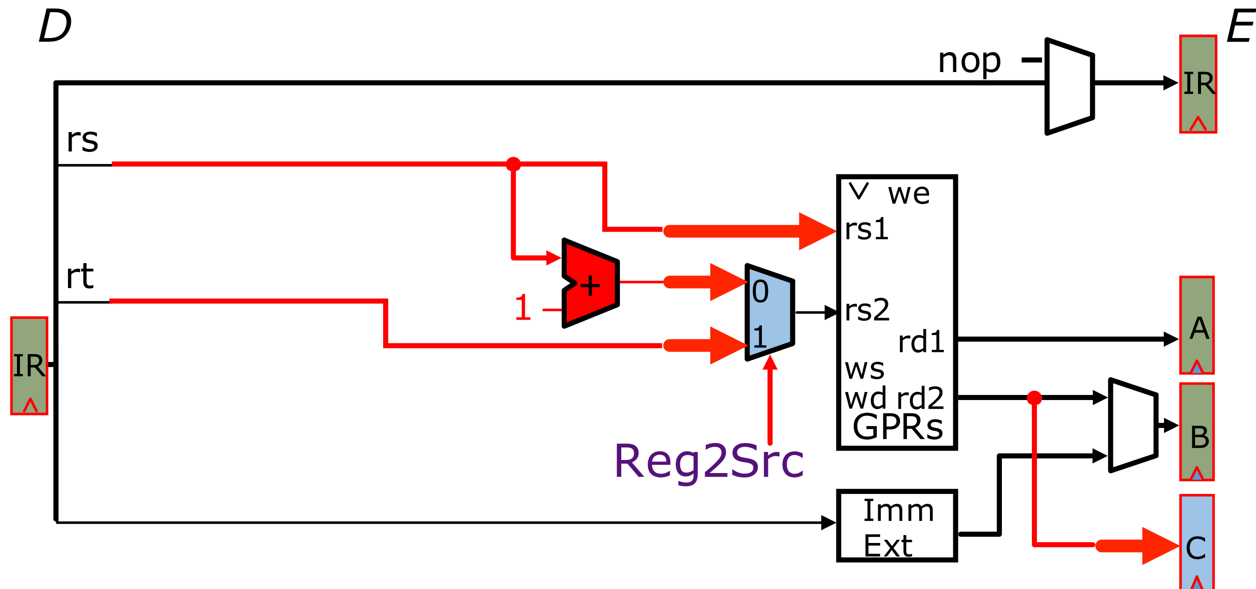

a) In the figure above, draw the datapath so that it handles BigMIPS loads (`LW`) and all other instructions except stores. Connect sources to the inputs of *rs1*, *C*, and the inputs to the mux before *rs2* (the four thick red arrows). You are free to add combinational logic as necessary.

The figure above shows logic for (*rs* + 1) but (*rs* OR 1) is also acceptable (and more efficient)! Reversing the inputs to the mux before rs2 is acceptable so long as Reg2Src is updated accordingly in Q2b.

b) Derive the select signal for the mux input to *rs2*, labeled as *Reg2Src*.

You are allowed to use any internal signals (e.g., `OpCode`, `IR`, `rd1`, etc.) but not other control signals (`ExtSel`, etc.).

It is acceptable to consider Stores as well, but not necessary, as indicated earlier.

Reg2Src = __**!(OpCode$_D$ == LW || OpCode$_D$ == SW)**____

It is acceptable to ignore Stores.

Reg2Src = ________**OpCode$_D$!= LW**_______

Question 3: Three Ports for Stores (5 points)

Alyssa's previous Decode stage implementation works for loads, but she realizes that stores must read three registers! (Two to form a 64-bit address and one for the data). As a first prototype, she adds a third register read port to the register file: *rs3* and *rd3*, as shown below.

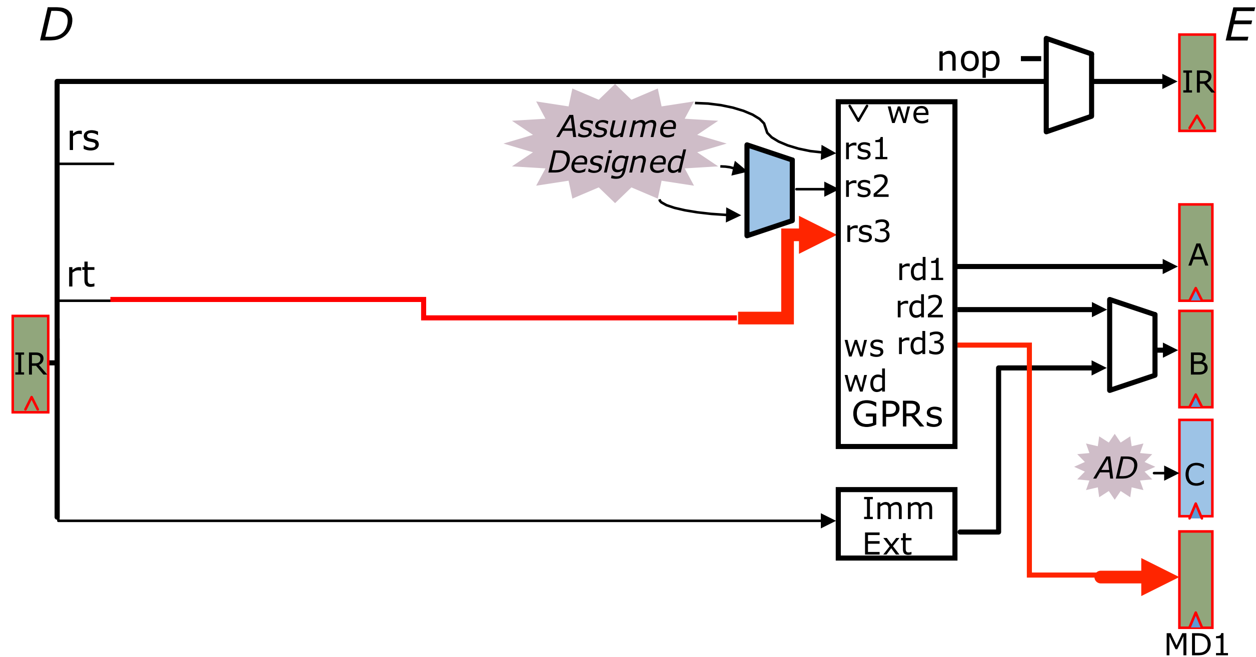

In the figure above, draw the completed datapath for a store. Assume the signals from Question 2 are already designed correctly to form a 64-bit effective address. Connect sources to the inputs of *rs3* and *MD1* (the two thick red arrows). You are free to add combinational logic as necessary.

The rt signal consists of a 5-bit register ID to retrieve data from the register file. It cannot be connected directly to the input of MD1 (which expects 32 bits of data).

Question 4: Three-Ported BigMIPS Resource Usage (8 Points)

Consider the following BigMIPS code sequence:

```
I1    LW    R7, 8(R2)
I2    LW    R5, 24(R6)
I3    ADD   R1, R5, R7
I4    SW    R1, 8(R8)
```

Complete the instruction flow diagram for this sequence of instructions. Assume used registers have been properly initialized. **Assume full bypassing** and correct stall logic to handle all features of BigMIPS. Use arrows to show forwarding of values from one stage to another. (In case you need it, page 21 has an extra/scratch instruction flow diagram.)

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
I1	F	D	E	M	W														
I2		F	D	D	E	M	W												
I3			F	F	D	D	E	M	W										
I4					F	F	D	E	M	W									
I5																			
I6																			

Alyssa finds that a third register read port is too expensive. Ben Bitdiddle suggests an ISA solution that doesn't need the third port. He replaces the Store Word instruction with two new instructions and adds a new register, *SMA*, that is part of architectural state. The new instructions and their formats are described in the table below.

Instruction	**Format and Description**

Store Word Address	`SWADDR offset(rs)`
	
	`SMA ← ((rs) << 32 | (rs+1)) + SignExt64(offset)`
	
	Calculate the 64-bit effective address as in the BigMIPS `LW` instruction. Load this address into register *SMA*. *rs* must be even.

Store Word Data	`SWDATA rt`
	
	`Mem[(SMA)] ← rt`
	
	Store the contents of register *rt* at the location addressed by *SMA*.

For Questions 5 and 6, assume that *Reg2Src* is updated appropriately for `SWADDR` and `SWDATA`.

The figure to the right shows Ben's changes to the Execute and Memory stages. The *MA* register of Figure C.1 is renamed to *LMA*. The *SMA* register is updated after the Execute stage of `SWADDR`. However, when `SWDATA` reaches the Memory stage, the *SMA* register must hold the address produced from the last `SWADDR` instruction. This behavior can be achieved through the Enable bit of the *SMA* register, *SMAen*.

With distinct memory address registers for loads and stores, Ben inserts a mux before the address field of the Data Memory, controlled by signal *MemAddrSrc*.

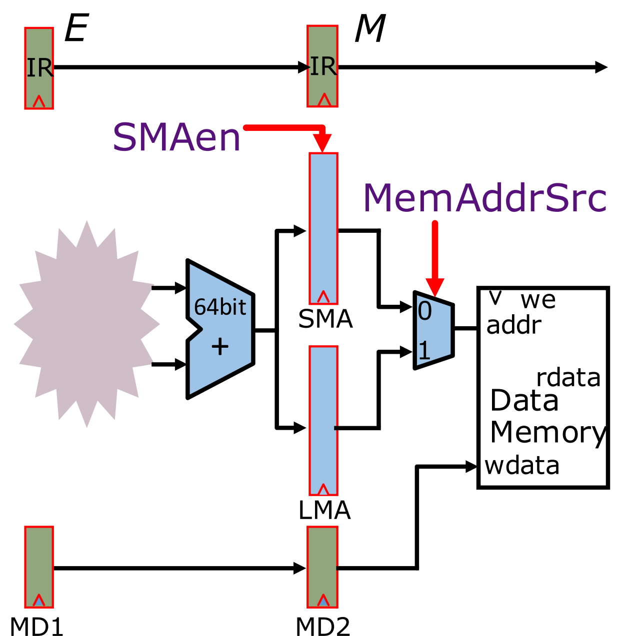

Question 5: Two Instructions for Stores (7 Points)

a) Derive the *SMAen* signal. Specify $\texttt{OpCode}_Z$ to indicate reading from the Z-stage *IR* register (e.g. $\texttt{OpCode}_D$ when reading the Decode stage *IR* register).

SMAen = __**$\texttt{OpCode}_E$ == SWADDR**____________

b) Derive the *MemAddrSrc* signal. Specify $\texttt{OpCode}_Z$ to indicate reading from the Z-stage *IR* register (e.g. $\texttt{OpCode}_D$ when reading the Decode stage *IR* register). Full credit will be given for use of "don't care" (*) when appropriate.

MemAddrSrc = Case *OpCode*__**M**__

____**SWDATA**________ => 0

____**LW**______________ => 1

___**otherwise**_______ => *

Question 6: Split-Store BigMIPS Resource Usage (8 Points)

We port the instruction sequence from Question 4 to use SWADDR and SWDATA:

```
I1    LW        R7, 8(R2)
I2    SWADDR    8(R8)
I3    LW        R5, 24(R6)
I4    ADD       R1, R5, R7
I5    SWDATA    R1
```

Complete the instruction flow diagram for this sequence of instructions. Assume used registers have been properly initialized. **Assume full bypassing** and correct stall logic to handle all features of BigMIPS. Use arrows to show forwarding of values from one stage to another. (In case you need it, page 21 has an extra/scratch instruction flow diagram.)

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
I1	F	D	E	M	W														
I2		F	D	E	M	W													
I3			F	D	E	M	W												
I4				F	D	D	E	M	W										
I5					F	F	D	E	M	W									
I6																			

Scratch Space

Use these extra pages if you run out of space or for your own personal notes. We will not grade this unless you tell us explicitly in the earlier pages.

Extra Instruction Flow Diagram

Use this as scratch space or if you need a new one to answer one of the questions in Part C.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
I1																			
I2																			
I3																			
I4																			
I5																			
I6																			
I7																			
I8																			