

Problem M3.2: Branch Prediction [? Hours]

This problem will investigate the effects of adding global history bits to a standard branch prediction mechanism. **In this problem assume that the MIPS ISA has no delay slots.**

Throughout this problem we will be working with the following program.

```
loop:
    LW    R4, 0(R3)
    ADDI  R3, R3, 4
    SUBI  R1, R1, 1
b1:
    BEQZ  R4, b2
    ADDI  R2, R2, 1
b2:
    BNEZ  R1, loop
```

Assume the initial value of R1 is n ($n > 0$).

Assume the initial value of R2 is 0 (R2 holds the result of the program).

Assume the initial value of R3 is p (a pointer to the beginning of an array of 32-bit integers).

All branch prediction schemes in this problem will be based on those covered in the lecture. We will be using a 2-bit predictor state machine, as shown below.




Figure M3.2-A: BP bits state diagram

In state 1X we will guess not taken. In state 0X we will guess taken.

Assume that b1 and b2 do not conflict in the BHT.

Problem M3.2.A

Program

What does the program compute? That is, what does R2 contain when we exit the loop?

Problem M3.2.B**2-bit branch prediction**

Now we will investigate how well our standard 2-bit branch predictor performs. Assume the inputs to the program are $n=8$ and $p[0] = 1, p[1] = 0, p[2] = 1, p[3] = 0, \dots$ etc. That is the array elements exhibit an alternating pattern of 1's and 0's. Fill out Table M3.2-1 (note that the first few lines are filled out for you). What is the number of mispredicts?

Table M3.2-1 contains an entry for every branch (either b1 or b2) that is executed. The Branch Predictor (BP) bits in the table are the bits from the BHT. For each branch, check the corresponding BP bits (indicated by the bold entries in the examples) to make a prediction, then update the BP bits in the following entry (indicated by the italic entries in the examples).

Problem M3.2.C**Branch prediction with one global history bit**

Now we add a global history bit to the branch predictor, as described in the lecture. Fill out Table M3.2-2, and again give the total number of mispredicts you get when running the program with the same inputs.

Problem M3.2.D**Branch prediction with two global history bits**

Now we add a second global history bit. Fill out Table M3.6-3. Again, compute the number of mispredicts you get for the same input.

Problem M3.2.E**Analysis I**

Compare your results from problems M3.2.B, M3.2.C and M3.2.D. When do most of the mispredicts occur in each case (at the beginning, periodically, at the end, etc.)? What does this tell you about global history bits in general? For a large n , what prediction scheme will work best? Explain briefly.

Problem M3.2.F**Analysis II**

The input we worked with in this problem is quite regular. How would you expect things to change if the inputs were random (each array element were equally probable to be 0 or 1). Of the three branch predictors we looked at in this problem, which one will perform best for this type of input? Is your answer the same for large and small n ?

What does this tell you about additional history bits: when are they useful and when do they hurt you?

System State		Branch Predictor		Branch Behavior	
PC	R3/R4	b1 bits	b2 bits	Predicted	Actual
b1	4/1	10	10	N	N
b2	4/1	10	10	N	T
b1	8/0	10	11	N	T
b2	8/0	11	11	N	T
b1	12/1	11	00		
b2	12/1				
b1					
b2					
b1					
b2					
b1					
b2					
b1					
b2					
b1					
b2					
b1					
b2					
b1					
b2					
b1					
b2					

Table M3.2-1

System State			Branch Predictor				Behavior	
PC	R3/R4	history bit	b1 bits		b2 bits		Predicted	Actual
			set 0	set 1	set 0	set 1		
b1	4/1	1	10	10	10	10	N	N
b2	4/1	0	10	10	10	10	N	T
b1	8/0	1	10	10	11	10		
b2	8/0							
b1	12/1							
b2	12/1							
b1								
b2								
b1								
b2								
b1								
b2								
b1								
b2								
b1								
b2								
b1								
b2								
b1								
b2								
b1								
b2								

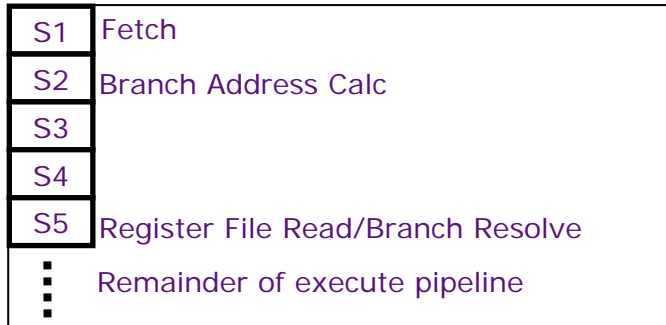
Table M3.2-2

System State			Branch Predictor								Behavior	
PC	R3/R4	history	b1 bits				b2 bits				Predicted	Actual
		bits	set 00	set 01	set 10	set 11	set 00	set 01	set 10	set 11		
b1	4/1	11	10	10	10	10	10	10	10	10	N	N
b2	4/1	01	10	10	10	10	10	10	10	10	N	T
b1	8/0	10	10	10	10	10	10	11	10	10		
b2	8/0											
b1	12/1											
b2	12/1											
b1												
b2												
b1												
b2												
b1												
b2												
b1												
b2												
b1												
b2												
b1												
b2												
b1												
b2												


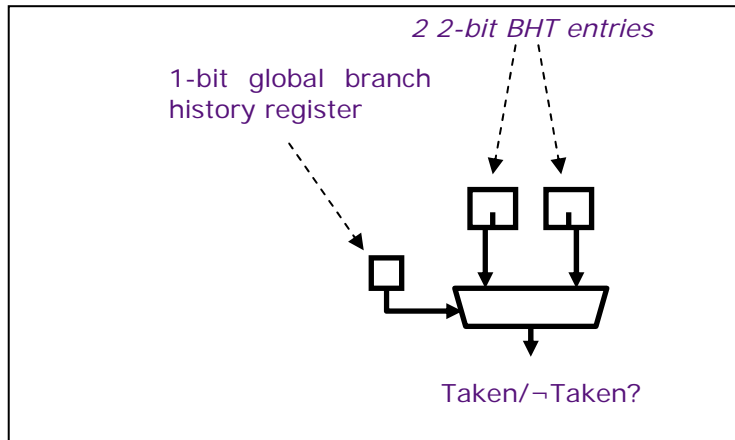
Table M3.2-3

Problem M3.3: Branch Prediction [? Hours]

Consider a CPU with a pipeline pictured on the right. The first stage of the pipeline fetches the instruction. The **second stage** of the pipeline recognizes branch instructions and **performs branch prediction** using a BHT. If the branch is predicted to be taken, it forwards the decoded target of the branch to the first stage, and kills the instruction in the first stage. The **fifth stage** of the pipeline reads the registers and **resolves the correct target of the branch**. If the branch target was mispredicted, the correct target is forwarded to the first stage, and all instructions in between are killed. The remaining stages finish the computation of the instruction.



The processor uses a **single global history bit** to remember whether the last branch was taken or not. There is only **one line in the BHT**, so the address of the branch instruction is not used for selecting the proper table entry. Each entry in the table is labeled as TW for Take Wrong, TR for Take Right, NTW for do Not Take Wrong and NTR for do Not Take Right, as pictured below. The setup of the BHT predictor is illustrated on the right.



In this question we will study execution of the following loop. This processor has **no** branch delay slots. You should assume that branch at address 1 is **never** taken, and that the branch at address 5 is **always** taken.

Instruction Label	Address	Instruction
LOOP	1	BEQ R2, R5, NEXT
	2	ADD R4, R4, 1
	3	MULT R3, R3, R4
NEXT	4	MULT R2, R2, 3847
	5	BNEZ R4, LOOP
	6	NOP
	7	NOP
	8	NOP
	9	NOP
	10	NOP

You should also **disregard any possible structural hazards**. The processor always runs at full speed, and there are **no pipeline bubbles** (except for those created by the branches).

Problem M3.3.A

Now we study how well the history bit works, when it is being **updated by the fifth stage** of the processor. The fifth stage also updates the BHT based on the result of a branch. The same BHT entry that was used to make the original prediction is updated.

Please fill in the table below.

You should fetch a new instruction every cycle. You should fill in the *Branch Prediction* and the *Prediction Correct?* columns for branch instructions only (note that the branch prediction actually happens one cycle **after** the instruction is fetched). You should fill in the *Branch Predictor State* columns whenever they are updated. Please **circle the instructions which will be committed**.

The first three committing instructions fetched have been filled in for you. You should enter enough instructions to **add 8 more committing instructions**. You may not need all the rows in the table.

Cycle	Instruction Fetched	Branch Prediction	Prediction Correct?	Branch Predictor State		
				Branch History	Last Branch Taken Predictor	Last Branch Not Taken Predictor
0	-	-		T	TW	TW
1	1	T	N			
2	2					
3	4					
4	5	T	Y			
5	6			NT	NTR	
6	2					
7	3					
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						
20						
21						
22						
23						
24						
25						
26						

Problem M3.3.B

Now we study how well the branch **history bit** works, when it **is being updated speculatively** by the second stage of the processor. If the branch is mispredicted, the fifth stage sets the branch history bit to the correct value. Finally, the fifth stage also updates the BHT based on the result of a branch. The same BHT entry that was used to make the original prediction is updated.

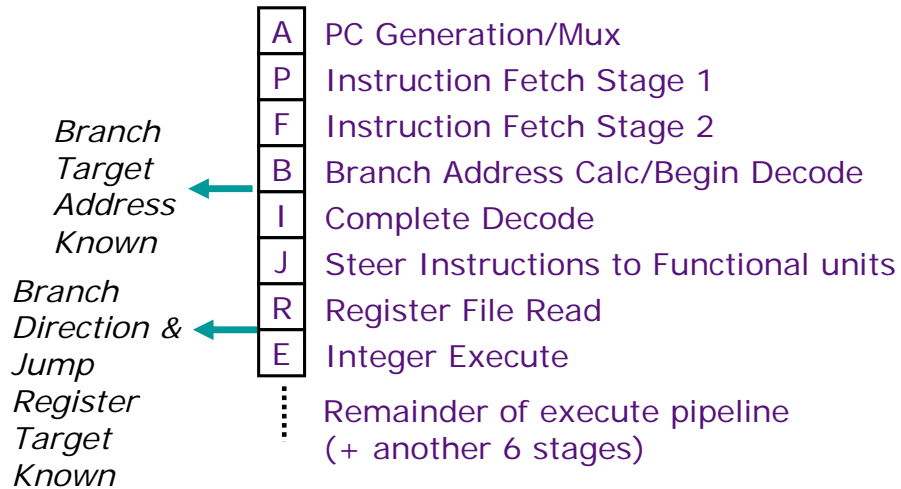
Please fill in the table below. The notation in the table should be same as in **M3.3.A**.

The first three committing instructions fetched have been filled in for you. You should enter enough instructions to **add 8 more committing instructions**. You may not need all the rows in the table.

Cycle	Instruction Fetched	Branch Prediction	Prediction Correct?	Branch Predictor State		
				Branch History	Last Branch Taken Predictor	Last Branch Not Taken Predictor
0	-	-		T	TW	TW
1	1	T	N			
2	2			T		
3	4					
4	5	T	Y			
5	6			NT	NTR	
6	2					
7	3					
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						
20						
21						
22						
23						
24						
25						
26						

Problem M3.4: Branch Prediction [? Hours]

Consider the fetch pipeline of UltraSparc-III processor. In this part, we evaluate the impact of branch prediction on the processor's performance. There are no branch delay slots.

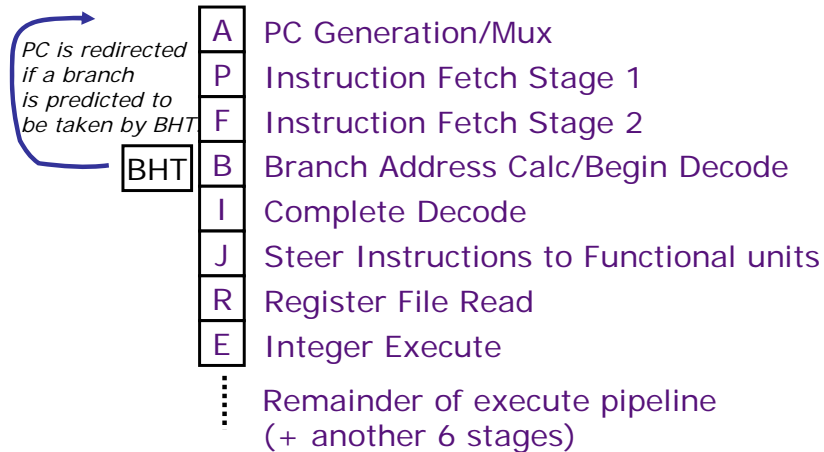


Here is a table to clarify when the direction and the target of a branch/jump is known.

Instruction	Taken known? (At the end of)	Target known? (At the end of)
BEQZ/BNEZ	R	B
J	B (always taken)	B
JR	B (always taken)	R

Problem M3.4.A

As a first step, we add a branch history table (BHT) in the fetch pipeline as shown on the next page. In the B stage (Branch Address Calc/Begin Decode), a conditional branch instruction (BEQZ/BNEZ) looks up the BHT, but an unconditional jump does not. If a branch is predicted to be taken, some of the instructions are flushed and the PC is redirected to the calculated branch target address. The instruction at PC+4 is fetched by default unless PC is redirected by an older instruction.



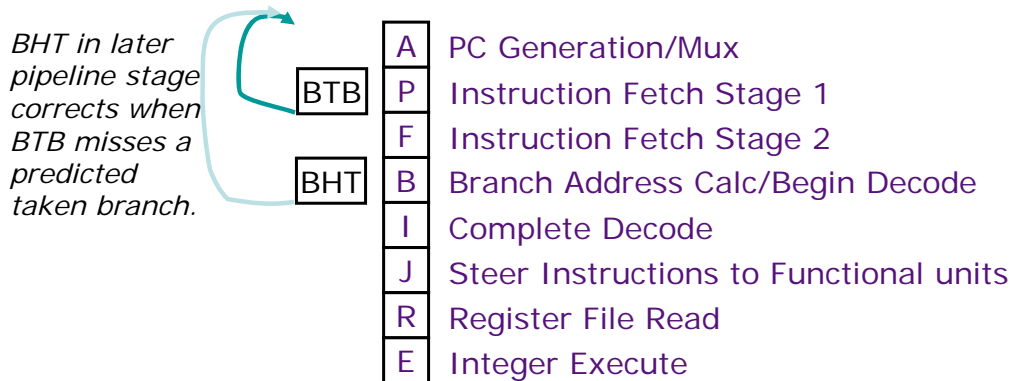
For each of the following cases, write down the number of pipeline bubbles caused by a branch or jump. If there is no bubble, you can simply put 0. (Y = yes, N= no)

	Predicted Taken?	Actually Taken?	Pipeline bubbles
BEQZ/ BNEZ	Y	Y	
	Y	N	
	N	Y	
	N	N	
J	Always taken (No lookup)	Y	
JR	Always taken (No lookup)	Y	

Problem M3.4.B

To improve the branch performance further, we decide to add a branch target buffer (BTB) as well. Here is a description for the operation of the BTB.

1. The BTB holds entry_PC, target_PC pairs for jumps and branches predicted to be taken. Assume that the target_PC predicted by the BTB is always correct for this question. (Yet the direction still might be wrong.)
2. The BTB is looked up every cycle. If there is a match with the current PC, PC is redirected to the target_PC predicted by the BTB (unless PC is redirected by an older instruction); if not, it is set to PC+4.



Fill out the following table of the number of pipeline bubbles (only for conditional branches).

	BTB Hit?	(BHT) Predicted Taken?	Actually Taken?	Pipeline bubbles
Conditional Branches	Y	Y	Y	
	Y	Y	N	
	Y	N	Y	Cannot occur
	Y	N	N	Cannot occur
	N	Y	Y	
	N	Y	N	
	N	N	Y	
	N	N	N	


Problem M3.4.C

We will be working on the following program:

ADDRESS	INSTRUCTION
0x1000	BR1: BEQZ R5, NEXT ; always taken
0x1004	ADDI R4, R4, #4
0x1008	MULT R3, R5, R3
0x100C	ST R3, 0(R4)
0x1010	SUBI R5, R5, #1
0x1014	NEXT: ADDI R1, R1, #1
0x1018	SLTI R2, R1, 100 ; repeat 100 times
0x101C	BR2: BNEZ R2, BR1
0x1020	NOP
0x1024	NOP
0x1028	NOP

Given a snapshot of the BTB and the BHT states on entry to the loop, fill in the timing diagram for one iteration (plus two instructions) on the next page. (Don't worry about the stages beyond the E stage.) We assume the following for this question.

- The initial values of R5 and R1 are zero, so BR1 is always taken.
- We disregard any possible structural hazards. There are no pipeline bubbles (except for those created by branches.)
- We fetch only one instruction per cycle.
- We use a two-bit predictor whose state diagram is shown below. In state 1X we will guess not taken; in state 0X we will guess taken. BR1 and BR2 do not conflict in the BHT.



- We use a two-entry fully-associative BTB with the LRU replacement policy.

Initial Snapshot

(Valid)	Entry	PC	Predicted Target PC
1	0x101C	0x1000	

BTB

	...
BR1	1 1
	...
BR2	0 0
	...

BHT

Address	Instruction	TIME →																						
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
0x1000	BEQZ R5, NEXT	A	P	F	B	I	J	R	E															
0x1014	ADDI R1, R1, #1																							
0x1018	SLTI R2, R1, 100																							
0x101C	BNEZ R2, LOOP																							
0x1000	BEQZ R5, NEXT																							
0x1014	ADDI R1, R1, #1																							

Timing diagram for M3.4.C

Problem M3.4.D

What will be the BTB and BHT states right after the 6 instructions in Question 9 have updated the branch predictors' states? Fill in (1) the BTB and (2) the entries corresponding to BR1 and BR2 in the BHT.

(Valid)	Entry PC	Predicted Target PC

BTB

	...
BR1	
	...
BR2	
	...

BHT