

**Computer Architecture
ECS 154B
Winter 2004**

Final Exam

Total: 100 points

Time: 120 minutes (1:30 PM - 3:30 PM)

Open book and open notes

Thu., March 18, 2004

Name:

Student Id:

Please SHOW ALL YOUR WORK if you want points!

1 Pipeline Hazards (2 points)

The following code contains a 'read after write data hazard that is resolved by forwarding:

```
add $2, $3, $4  
add $5, $2, $6
```

Consider the similar situation in which a memory read occurs after a memory write:

```
sw $7, 100($2)  
lw $8, 100($2)
```

1. How do the above two situations differ? (1 point)

2. Give a scheme for resolving the potential "read after write" problem in second situation? (1 point)

2 Design Forwarding Unit - Check the Handout/Book (13 points)

The forwarding unit could be moved to the ID stage and forwarding decisions could be made earlier. The results of these decisions would need to be passed along with the instruction and used in the EX stage when actual forwarding would take place. This modification would speed up the EX stage and might allow for possible cycle-time improvement.

1. Draw the datapath similar to one shown on page 484 given to you as a handout and show the modification you would make. (5 points)

2. What is the change in the ID/EX register? (2 points)

3. Give new forwarding equations for EX hazards and MEM hazards similar to ones appearing on pages 480-483 (See the handout). (3+3 points)

3 Branch Predictors (15 points)

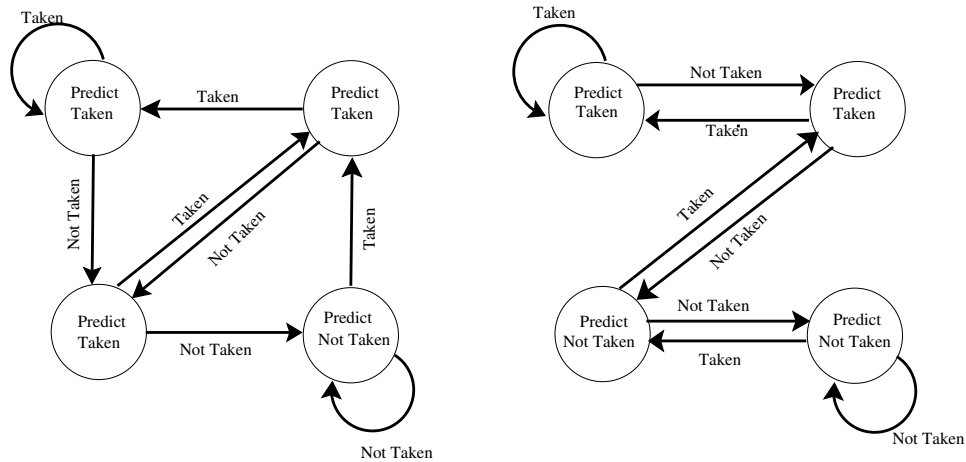


Figure 1: State diagram of two branch predictors - Predictor 1 and Predictor 2.

Figure 1 shows the state diagram of two branch predictors referred to as Predictor 1 and Predictor 2 respectively. We will also use Predictor 3 which is a two bit predictor discussed in the class. Please answer the questions below to compare the predictors: Assume all branch predictors begin with the prediction TAKEN.

1. Give a sequence of actual branch behavior for which Predictor 1 has higher accuracy than Predictor 2? (2 points)

2. Give a sequence of actual branch behavior for which Predictor 2 has higher accuracy than Predictor 1? (2 points)

3. Give a sequence of actual branch behavior for which Predictor 1 has higher accuracy than Predictor 3? (2 points)

4. Give a sequence of actual branch behavior for which Predictor 3 has higher accuracy than Predictor 1? (2 points)

4 Virtual Memory (15 points)

A paged memory with a one-level page table has the following parameters: The pages are 2^P bytes long; virtual addresses are V bits long.

The page-table starts at physical address $PTBL$; and each page-table entry is a 4-byte longword, so that, given a virtual address, the relevant page-table entry can be found at $PTBL + (\text{page number}) * 4$. Answer the following in terms of the parameters P and V :

1. How many bits long is the "offset in page" field? (2 points)
2. How many bits long is the "virtual page number" field? (2 points)
3. How many entries does the page table have, and what is the highest address occupied by a page-table entry? (3 points)
4. How many pages long is the page table? (2 points)
5. What is the smallest value of P such that the page table fits into one page? (2 points)
6. The following table shows the first 8 entries in the page map. Recall that the valid bit is 1 if the page is resident in physical memory and 0 if the page is on disk or hasn't been allocated.

Virtual Page	Valid bit	Physical Page
0	0	7
1	1	9
2	0	3
3	1	2
4	1	5
5	0	5
6	0	4
7	1	1

If there are $1024(2^{10})$ bytes per page, what is the physical address corresponding to the decimal virtual address 3956? (4 points)

5 Parallel Programming (15 points)

The following FORTRAN program is to be executed on a computer, and a parallel version is to be executed on a 32-computer cluster.

```
L1:      DO 10 I = 1, 1024
L2:          SUM(I) = 0
L3:          DO 20 J = 1, I
L4: 20      SUM (I) = SUM(I) + I
L5: 10  CONTINUE
```

Suppose lines 2 and 4 each take two machine cycle times, including all processor and memory-access activities. Ignore the overhead caused by the software loop control statements (lines 1, 3, 5) and all other system overhead and resource conflicts.

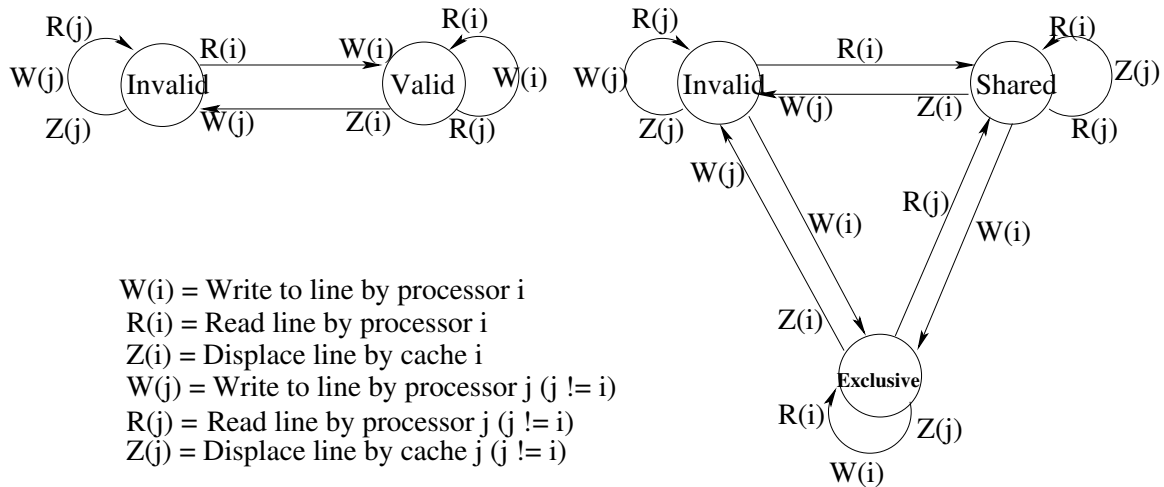
(a.) What is the total execution time (in machine cycle times) of the program on a single computer? (2 points)

(b.) Divide the I-loop iterations among the 32 computers as follows: Computer 1 executes the first 32 iterations (I=1 to 32), processor 2 executes the next 32 iterations, and so on. What are the execution time and speedup factor compared with part (a)? (Note that the computational workload, dictated by the J-loop, is unbalanced among the computers.) (4 points)

(c.) Explain how to modify the parallelizing to facilitate a balanced parallel execution of all the computational workload over 32 computers. By a balanced load is meant an equal number of additions assigned to each computer with respect to both loops. (7 points)

(d.) What is the minimum execution time resulting from the parallel execution on 32 computers? What is the resulting speedup over a single computer? (2 points)

6 Cache Coherence (20 points)



Note: State diagrams are for a given line in cache i (initiating processor)

Figure: Two cache coherence Protocols. We will call them protocols 1 and 2 respective

Figure 2: Cache Coherence Protocols 1 and 2.

Figure 2 shows the state diagram of two cache coherence protocols. We will refer to the protocol with two states as Protocol 1 and the one with three states as Protocol 2. Both protocols are *write invalidate snooping protocols*.

Please answer the following questions. *You may use the terms 'cache line fill' and 'dirty line copyback' to explain your answers.*

For Protocol 1:

1. Provide sequence of actions when a processor reads a line present in its cache? (1 point)

2. Provide sequence of actions when a processor wants to read a line not present in its cache? (2 points)

3. Provide sequence of actions when a processor writes to a line present in its cache? (1 point)

11. Provide sequence of actions when a processor wants to write to a line not present in its cache? (2 points)

12. Provide sequence of actions when a processor 'displaces' a line from its cache? (1 point)

13. Illustrate by a clear example, an advantage of MESI protocol over Protocol 2. (2 points)

14. Give an advantage of Protocol 2 over MESI protocol. (1 point).

7 Interconnection Networks (20 points)




Figure 3: A 4X4 mesh of trees.

A $N \times N$ mesh of trees is a $N \times N$ mesh where every row and every column also form the leaves for a binary tree. That is, there is a binary tree attached to every row and every column. This is an indirect network - processors are only attached to the switches in the mesh, not in the switches in the trees. For example, in Fig. 3, 16 processors will be attached to the 16 switches which are interconnected by solid lines.

1. How many switches does the $N \times N$ mesh of trees have? (2 points)
2. What is the diameter of the $N \times N$ mesh of trees? (4 points)
3. What is the total bandwidth of the $N \times N$ mesh of trees? (4 points)

4. What is the bisection bandwidth of the $N \times N$ mesh of trees? (4 points)

5. If a 2×2 crossbar switch costs \$5 and cost is proportional to the square of the number of inputs, what is the total cost the switches in a 16×16 mesh of trees network? (6 points)