CE1-R3: ADVANCED COMPUTER ARCHITECTURE

NOTE:

1. **Answer question 1 and any FOUR questions from 2 to 7.**
2. **Parts of the same question should be answered together and in the same sequence.**

Time: 3 Hours **Total Marks: 100**

1.

a) Discuss two areas of computer design where pipelining is appropriate.
b) Give (at least four) major characteristics of a systolic array processor.
c) What are the possible ways of organizing the operating system of a multiprocessor to meet its objectives?
d) Consider the following FORTRAN DO loop:

```
        SUM=0
        DO 10 I=1, 1000, 1
10      SUM = SUM + B(I)
```

Perform some modification to achieve speedup.
e) What are different types of pipeline hazards?
f) What are the factors to be considered in deciding cache size?
g) Outline some specific properties of RISC systems.

(7x4)

2.

a) Explain the concept of pipelining in superscalar processors.
b) What is dynamic branch-prediction? Discuss one dynamic branch prediction scheme.
c) What are branch-target buffers? Explain with an example.
d) Determine the total branch penalty for a branch-target buffer assuming the penalty cycles for individuals mispredictions from the following table.

Instruction in buffer	Prediction	Actual branch	Penalty cycles
yes	taken	taken	0
yes	taken	not taken	2
No		taken	2
No		not taken	0

Make the following assumptions about the predictions about the prediction accuracy and hit rate:

- Prediction accuracy is 90% (for instructions in the buffer).
- Hit rate in the buffer is 90% (for branches predicted taken).

Assume that 60% of the branches are taken.

(3+5+4+6)

3.

a) Discuss different software approaches used to exploit instruction level parallelism.
b) Consider a loop like the one below:

```
for (i=1; i<=100; i=i+1) {
     A[i+1] = A[i] + C[i];       /*S1*/
     B[i+1] = B[i] + A[i+1];     /*S2*/
```

Assume that A,B, and C are distinct, nonoverlapping arrays. What are the data dependences among the statements S1 and S2 in the loop?

c) Consider a loop like this one:

```
for (i = 1; i<=100; i=i+1) {
     A[i] = A[i] + B[i];         /*S1*/
     B[i+1] = C[i] + D[i];       /*S2*/
}
```

What are the dependences between S1 and S2? Is this loop parallel? If not, Show how to make it parallel.

d) The following loop has multiple types of dependences. Find all the true dependences, output dependences, and antidependences, and eliminate the output dependences and antidependences by renaming.

```
for   (i = 1; i<=100; i = i+1)    {
      Y[i]  = X[i] / c;      /*S1*/
      X[i]  = X[i] +c; /*S2*/
      Z[i]  = Y[i] +c; /*S3*/
      Y[i]  = c - Y[i];      /*S4*/
}
```

(3+4+5+6)

4.

a) Assume we have a computer where the clock per instruction (CPI) is 1.0 when all memory accesses hit in the cache. The only data accesses are loads and stores, and these total 50% of the instructions. If the miss penalty is 25 clock cycles and the miss rate is 2%, how much faster would the computer be if all instructions were cache hits?

b) A pipeline *P* is found to provide a speedup of 6.16 when operating at 100 MHz and an efficiency of 88 percent.

i) How many stages does *P* have?

ii) What are *P*'s MIPS and CPI performance levels?

c) Let's use an in-order execution computer for the first example, such as the Ultra-SPARC III. Assume the cache miss penalty is 100 clock cycles, and all instructions normally take 1.0-clock cycles (ignoring memory stalls). Assume the average miss rate is 2%, there is an average of 1.5 memory references per instruction, and the average number of cache misses per 1000 instructions is 30. What is the impact on performance when behavior of the cache is included? Calculate the impact using both misses per instruction and miss rate.

(4+8+6)

5.

a) What is a nonblocking cache? Discuss its advantages.

b) Explain compiler-controlled prefetching technique.

c) What is a virtual cache? Explain why virtual caches are not popular?

d) What can interleaving and wide memory buy? Consider the following description of a computer and its cache performance:

Block size = 1 word
Memory bus width = 1 word
Miss rate = 3%
Memory accesses per instruction = 1.2
Cache miss penalty = 64 cycles (as above)
Average cycles per instruction (ignoring cache misses) = 2

If we change the block size to 2 words, the miss rate falls to 2%, and a 4-word block has a miss rate of 1.2%. What is the improvement in performance of interleaving two ways versus doubling the width of memory and the bus?

(4+4+5+5)

6.

a) Discuss Flynn's classification of processors.

b) Compare shared memory multiprocessor architecture and distributed memory architecture.

c) Suppose you want to achieve a speedup of 80 with 100 processors. What fraction of the original computation can be sequential?

d) Suppose we have an application running on a 32-bit multiprocessor, which has a 400 ns time to handle reference to a remote memory. For this application, assume that all the references except those involving communication hit in the local memory hierarchy, which is slightly optimistic. Processors are stalled on a remote request, and the processor clock rate is 1 GHz. If the base IPC (assuming that all references hit in the cache) is 2, how much faster is the multiprocessor if there is no communication versus if 0.2% of the instructions involve a remote communication reference?

(3+4+6+5)

7.

a) Compare shared and switched interconnection media.

b) Describe the following terminologies associated with multiprocessor operating systems and MIMD algorithms:
   i) Protection mechanisms
   ii) Scheduling
   iii) Degree of decomposition of a parallel algorithm.

c) The CM-5 supercomputer used wormwhole routing, with each switch buffer being just 4 bits per port. Compare efficiency of store-and-forward versus wormhole routing for a 128-node machine using a CM-5 interconnection sending a 16-byte payload. Assume each switch takes 0.25 μs and that the transfer rate is 20MB/sec.

(4+[3×3]+5)