Question 1 (5 points)

Consider a cache with the following specifications
Address space is 1024 words.
The memory is word addressable
The size of the cache is 8 blocks; each block is 4 words (32 words cache).
The cache is 2-way associative
Each block in the cache has a valid bit associated with it.
What is the content of the cache after the following sequence of memory references? (Show the cache contents in the form of a table with "at least" the following columns: valid, tag,)
The reference pattern is (in hex)
36A, 36B, 26E 218,173,170, 169 with the first reference as 36A

<u>Question 2</u> (5 points 2-3)

Assume that the following is the branch outcome of a single branch (T means taken N meant not taken)

T T N N T N N N T T N T N T

If we assume a single bit branch prediction, and assume that this is the only branch in the program. The branch prediction table was initialized to 0 (not taken)

How many misprediction?

A table like that may be helpful

Prediction before branch	**Branch outcome**	**Misprediction? 1 is yes, 0 no**
N	T	
	T	
	N	
	N	
	T	
	N	
	N	
	N	
	T	
	T	
	N	
	T	
	N	
	T	

If we use a 2-bit saturating counter for branch prediction, and we assume that it was initialized to 01, how many branch misprediction? (00 means deep not taken, 01 means borderline not taken, 10 means borderline taken, 11 means deep taken).

Prediction before branch	**Branch outcome**	**Misprediction? 1 is yes, 0 no**
01	T	
	T	
	N	
	N	
	T	
	N	
	N	
	N	
	T	
	T	
	N	
	T	
	N	
	T	

Question 3 (7 points 2-2-3)

Consider a transaction processing-based computer system. The specifications are 25MIPS CPU costing $40,000 and the I/O require 250,000 instructions. You are required to choose the disks and you have 2 choices.

1. 300MB disks each one can perform 30 IOPS (I/O Per Second) and costs $8,000.
2. 150MB disks each can perform 40 IOPS and costs $6,000

All disks must be of the same type, and consider only the CPU and disks as potential bottlenecks (i.e. ignore the bus and controllers)

a) Find the maximum transaction rate in IOPS the system can achieve
b) Find the cost of IOPS for disks of types 1 and 2 to achieve the max. rate
c) Redo (b) if we put a condition that the disk space should be at least five times the transaction rate multiplied by 10^7 bytes.

Question 4 (6 points)

Consider a single issue Tomasulo's-based CPU with the following specifications.

- Load and stores take 2 cycles, one cycle to calculate the effective address using the INT unit, the second is to access the memory.
- FP add takes 4 cycles, FP multiply takes 8 cycles.
- There is one INT ALU, one FP adder, and one FP multiplier each with a single reservation station.
- Unlimited load and store buffers.
- CDB contention is resolved in favor of the instruction issued first.
- The reservation unit is considered free after the write back stage.

Consider the following piece of code. Fill in the table where Issued means the clock cycle the instruction was issued in. Executed is the range of cycles the instruction is being executed (for example 2-4 means the instruction will be in the execution stage for cycles 2,3, and 4). Finally write back is the cycle during which the result is written back

```
LD      F2, 0(R1)
MULD    F4, F2, F0
LD      F6, 0(R2)
ADDD    F6, F4, F6
SDD     0(R2), F6
ADDI    R1, R1, #8
ADDI    R2, R2, #8
SUBI    R3, R1, #800
```

Instruction	Issued	Executed	Write-back
LD			
MULD			
LD			
ADDD			
SDD			
ADDI			
ADDI			
SUBI			

Question 5 (6 points)

Consider the following code that adds a constant to an array

```
Loop:   LD      F0,0(R1)
        ADD     F4,F0,F2
        SD      0(R1),F4
        ADDI    R1,R1, #-8
        BNE     R1,R2,Loop
```

Show how that code will be scheduled on a VLIW that could issue 2 memory references, one FP operation, and one integer operation (BNE is done using the int unit).
Consider the following set of latencies

- 1 branch delay slot
- 1 cycle latency from load to FP or INT operation
- 1 cycle latency from the any function unit to store

Unroll the loop twice only

Cycles	Memory	Memory	FP Operation	INT Operation
1				
2				
3				
4				
5				
6				
7				
8				

Question 6 (4 points 1-1-1-1)

A You have a 2-way set associative cache with FIFO replacement policy and you noticed that the miss rate is too high. Which of the following suggestions can improve the miss rate? Circle the correct one(s)

(a) Change the replacement policy from FIFO to LRU
(b) Increase the size of the cache
(c) In crease the size of the page table
(d) Change it to direct mapped cache

B A byte addressable processor has a physical address space of 36 bits and the word consists of 4 bytes. The cache size is 128KB with a block size of 32 bytes. The cache is a direct mapped cache.
What is the number of bits in the three fields of the physical address? (offset, index, and tag)

C In a cache increasing the line size results in decreasing the capacity miss without changing the associativity or the size of the cache

TRUE **FALSE**

D Doubling the cache associativity can decrease the conflict miss without changing the cache size or the line size

TRUE **FALSE**

Question 7 (7 Points)

Assume that the CPI for a certain CPU with a perfect cache is 2.0. and also assume 1.33 memory reference per instructions. A cache hit is done in 1 cycle, the cache miss takes 4+2* number of words/block cycles. You have a choice of 2 cache organization.
Direct mapped cache with a block size of 16 words, and a miss rate of 1.8%
Direct mapped cache with a block size of 32 words and a miss rate of 1%

What is the CPI under these two cache organizations?