NOTE: If some questions are unclear or not well defined to you, you can make your own assumptions and state them clearly in the answer sheet.

1. (14% total) The following table shows the percentage of MIPS instructions executed by category for average of SPEC2000 integer programs and SPEC2000 floating point programs.

Instruction class	MIPS examples	Average CPI	Frequency	
			Integer	Floating point
Arithmetic	add, sub, addi	1.0 clock cycles	24%	48%
Data transfer	lw, sw, lb, sb, lui	1.4 clock cycles	36%	39%
Logical	and, or, nor, andi, ori	1.0 clock cycles	18%	5%
Conditional branch	beq, bne, slt, slti	1.7 clock cycles	18%	6%
Jump	j, jr, jal	1.2 clock cycles	4%	2%

1.1 (4%) Using the average instruction mix information for the program SPEC2000fp, find the percentage of all memory accesses (both data and instruction) that are for reads. Assume that two-thirds of data transfers are loads.

1.2 (4%) Compute the effective CPI for MIPS. Average the instruction frequencies for SPEC2000int and SPEC2000fp to obtain the instruction mix.

1.3 (6%) Consider an architecture that is similar to MIPS except that it supports update addressing for data transfer instructions. If we run SPEC2000int using this architecture, some percentage of the data transfer instructions will be able to make use of the new instructions, and for each instruction changed, one arithmetic instruction can be eliminated. If 25% of the data transfer instructions can be changed, which will be faster for SPEC2000int, the modified MIPS architecture or the unmodified architecture? How much faster? (Assume that both architectures have CPI values as given in the above table and that the modified architecture has its cycle time increased by 20% in order to accommodate the new instructions.)

2. (12% total) Consider two different implementations, I1 and I2, of the same instruction set. There are three classes of instructions (A, B, and C) in the instruction set. I1 has a clock rate of 4 GHz, and I2 has a clock rate of 2 GHz. The following table shows the average number of cycles for each instruction class on I1 and I2.

Class	CPI on I1	CPI on I2	C1 Usage	C2 Usage	C3 Usage
A	2	1	40%	40%	50%
B	3	2	40%	20%	25%
C	5	2	20%	40%	25%

The table also contains a summary of average proportion of instruction classes generated by three different compilers. C1 is a compiler produced by the makers of I1, C2 is produced by the makers of I2, and C3 is a third-party product. Assume that each compiler uses the same number of instructions for a given program but that the instruction mix is as described in the table.

2.1 (6%) Using C1 on both I1 and I2, how much faster can the makers of I1 claim I1 is compared to I2?

2.2 (6%) Which computer and compiler would you purchase if all other criteria are identical, including cost?

3. (10% total) You are going to enhance a computer, and there are two possible improvements: either make multiply instructions run four times faster than before, or make memory access instructions run two times faster than before. You repeatedly run a program that takes 100 seconds to execute. Of this time, 20% is used for multiplication, 50% for memory access instructions, and 30% for other tasks.

3.1 (4%) What will the speedup be if both improvements are made?

3.2 (6%) You are going to change the program described in Problem 3 so that the percentages are not 20%, 50%, and 30% anymore. Assuming that none of the new percentages is 0, what sort of program would result in a tie with regard to speedup (i.e., the same speedup) between the two individual improvements? Provide both a formula and some examples.

4. (16% total) Figure 1 gives the datapath of a pipelined processor with forwarding and hazard detection.

4.1 (3%) We have a program of 1000 instructions in the format of "lw, add, lw, add, ..." The add instruction depends (and only depends) on the lw instruction right before it. The lw instruction also depends (and only depends) on the add instruction right before it. If the program is executed on the pipelined datapath of Figure 1, what would be the actual CPI?

4.2 (3%) What would be the actual CPI for the program in Problem 4.1 without forwarding?

4.3 (4%) Consider executing the following code on the pipelined datapath of Figure 1. How many cycles will it take to execute this code?

```
lw      $4, 100($2)
sub     $6, $4, $3
add     $2, $3, $6
add     $7, $5, $2
```

4.4 (6%) With regard to the code in Problem 4.3, explain what the forwarding unit is doing during the sixth cycle of execution. If any comparisons are being made, mention them.

5. (16% total) We have a program core consisting of several conditional branches. The program core will be executed thousands of times. Below are the outcomes of each branch for one execution of the program core (T for taken, N for not taken).

　　Branch 1: T-T-T-T,　Branch 2: N-N-N-N-N,　Branch 3: T-N-T-N-T-N,　Branch 4: T-T-T-N-T

Assume the behavior of each branch remains the same for each program core execution. For dynamic schemes, assume each branch has its own prediction buffer and each buffer initialized to the same state before each execution. List the predictions and calculate the prediction accuracy for the following branch prediction schemes:

5.1 (5%) Always-taken

5.2 (5%) 1-bit predictor (initialized to predict taken)

5.3 (6%) 2-bit predictor as shown in Figure 2 (initialized to weakly predict taken)

6. (12% total) The average memory access time (AMAT) is possibly useful as a figure of merit for different cache systems.

6.1 (3%) Find the AMAT for a processor with a 2 ns clock, a miss penalty of 20 clock cycles, a miss rate of 0.06 misses per instruction, and a cache access time (including hit detection) of 1 clock cycle. Assume that the read and write miss penalties are the same and ignore other write stalls.

6.2 (3%) Suppose we can improve the miss rate to 0.04 misses per reference by doubling the cache size. This causes the cache access time to increase to 1.2 clock cycles. Using the AMAT as a metric, determine if this is a good trade-off.

6.3 (6%) If the cache access time determines the processor's clock cycle time, which is often the case, AMAT may not correctly indicate whether one cache organization is better than another. If the processor's clock cycle time must be changed to match that of a cache, is this a good trade-off? Assume the processors are identical except for the clock rate and the number of cache miss cycles; assume 1.5 references per instruction and a CPI without cache misses of 2. The miss penalty is 20 cycles for both processors.

7. (20% total) Consider three processors with different cache configurations:
　　Cache 1: Direct-mapped with two-word blocks
　　Cache 2: Direct-mapped with four-word blocks
　　Cache 3: Two-way set associative with four-word blocks
　The following miss rate measurements have been made:
　　Cache 1: Instruction miss rate is 3.75%; data miss rate is 5%
　　Cache 2: Instruction miss rate is 2%; data miss rate is 4%
　　Cache 3: Instruction miss rate is 2%; data miss rate is 3%

For these processors, one-half of the instructions contain a data reference. Assume that the cache miss penalty is 6 + Block size in words. The CPI for this workload was measured on a processor with *Cache 1* and was found to be 2.0.

7.1 (6%) Assuming a cache of 16K blocks and a 32-bit address, find the total number of sets and the total number of tag bits for *Cache 1*, *Cache 2*, and *Cache 3*.

7.2 (6%) Determine which processor spends the most cycles on cache misses.

7.3 (8%) The cycle times for the processors in Problem 7.2 are 400 ps for the first and second processors and 300 ps for the third processor. Determine which processor is the fastest and which is the slowest.
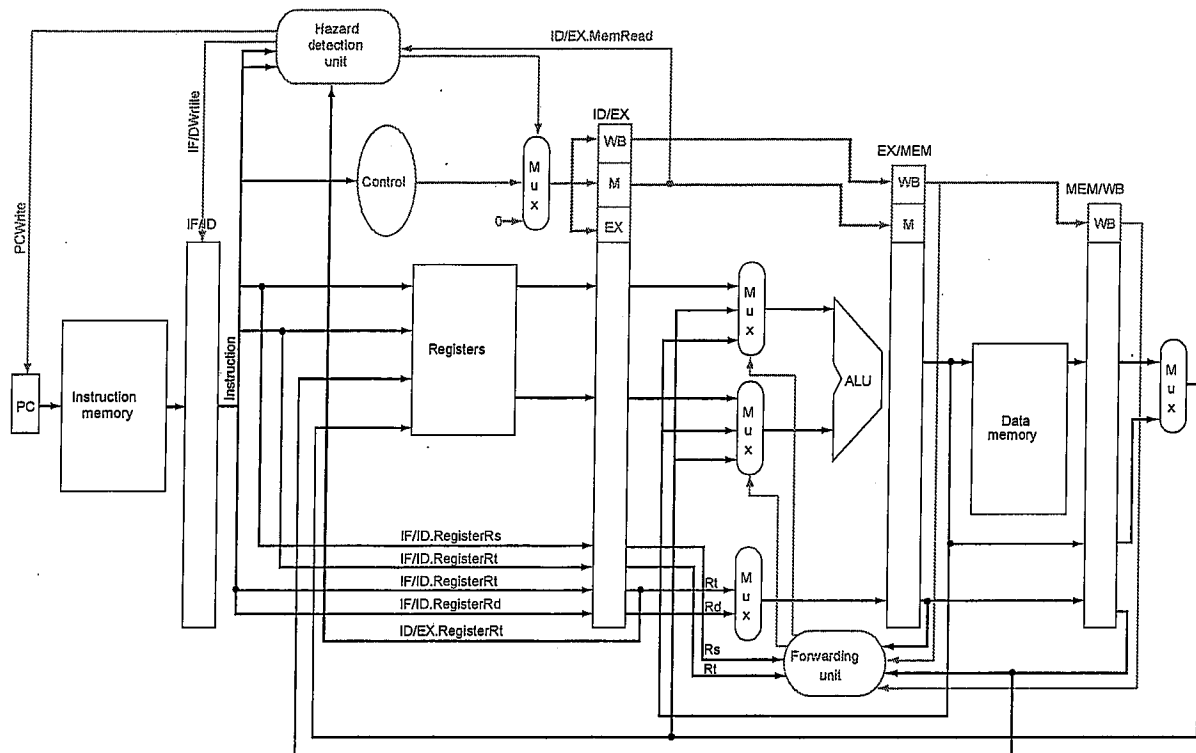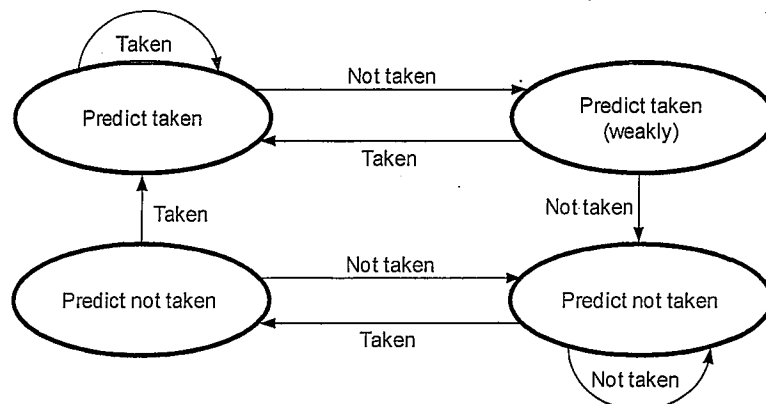

Figure 1

Figure 2

1. [Sorting: 10%]

Please prove that any comparison-based sorting algorithm requires O(*n lgn*) computation time in the worst case, where *n* is the number of objects to be sorted.

2. [Basic Data Structures: 10%]

(1) From the viewpoint of an operating system, please give practical applications of using a stack, queue, and linked list. (3%)

(2) Given a prefix expression "+ + (2) × (3) (4) × + × (5) (6) (7) (8)", please convert it to the infix expression and calculate the result. (2%)

(3) What is the assumption of simple uniform hashing? Under this assumption, given a hash table with *m* slots that stores *n* elements, what is the expected time taken by an unsuccessful search when the hash collisions are resolved by chaining? (3%)

(4) Please explain the property of a binary search tree. (2%)

3. [Graph: 15%]

You are given the adjacency matrix of an undirected graph:

	A	B	C	D	E	F	S
A	0	4	7	0	0	0	0
B	4	0	9	8	0	0	0
C	7	9	0	11	6	1	0
D	0	8	11	0	0	10	0
E	0	0	6	0	0	3	5
F	0	0	1	10	3	0	2
S	0	0	0	0	5	2	0

Starting from node S, please answer the following questions.

(1) Draw the corresponding graph. (3%)

(2) Write down the sequence of nodes that you visit by using breadth-first search. If there are multiple choices to select nodes, just follow the alphabetical order. (4%)

(3) Write down the sequence of nodes that you visit by using depth-first search. If there are multiple choices to select nodes, just follow the alphabetical order. (4%)

(4) Show the sequence of edges added to the minimum spanning tree based on the Prim's algorithm. Please use the edge weight to represent each edge in your answer. (4%)

4. [Advanced Data Structures: 15%]

(1) Please give the five properties of a binomial tree. (5%)

(2) Consider a Fibonacci heap as follows:

The number in each node indicates its key value. Suppose that we decrease key 48 to 17 and then decrease key 37 to 7. Please show the final result of the revised heap. (5%)

(3) Why does a B-tree can help reduce the cost of disk access? (5%)

5. **[Process: 10%]**
   (1) Please explain the five states of a process. (5%)
   (2) What is the risk of terminating a single thread by calling the exit() procedure in a multi-threading program? (2%)
   (3) Except for CPU utilization and throughput, please explain the three criteria to measure the performance of a process scheduling algorithm. (3%)

6. **[File System and Disk: 10%]**
   (1) Except for single-level and two-level directories, what are the three popular schemes to define the logical structure of directories? (3%)
   (2) Suppose that the range of a disk's cylinders is [1, 200] and the disk head currently stops at cylinder 55. Let the disk queue contain the requests of blocks on cylinders {99, 187, 39, 124, 14, 131, 65, 70}. Please show how the disk head moves in the SSTF, SCAN, C-SCAN, and C-LOOK disk scheduling schemes. (6%)
   (3) What is a component unit of the distributed file system? (1%)

7. **[Memory: 15%]**
   (1) Let the time to search the TLB and to access memory is 25 ns and 95 ns, respectively. Supposing that the TLB has 80% hit ratio, what is the effective memory-access time? Notice that you have to give the calculation. (3%)
   (2) Please explain both external and internal fragmentation in memory allocation. (4%)
   (3) From the viewpoint of a process, what is thrashing? (2%)
   (4) What is memory-mapped I/O? Why do we need it? (3%)
   (5) Why can virtual memory speed up process creation? (3%)

8. **[Synchronization: 15%]**
   (1) Please explain the three synchronization mechanisms for threads to access a shared data item in UNIX systems. (6%)
   (2) Suppose that three processes P1, P2, and P3 respectively have timestamps 17, 25, and 30, and P2 currently holds a resource requested by both P1 and P3. Please give the results of wait-die and wound-wait schemes used to handle deadlocks. (4%)
   (3) Please explain how the two-phase locking protocol works to check the correctness of a non-serial schedule for atomic transactions. (5%)

There are 7 problems in this test. Note that you should write down detailed steps for the solution to each problem; otherwise, no credits for that problem will be given.

1. Explain the following terminologies.
   (a) [3%] Partial ordering relation
   (b) [3%] Planar graph
   (c) [4%] Left coset

2. Let $A = \{a, b, c, d, e, f, g, h\}$ and R is the equivalence relation on A.
   (a) [6%] If R induces the partition $A = \{a\} \cup \{b, c\} \cup \{d, e\} \cup \{f, g, h\}$, what is R?
   (b) [6%] If R induces the partition $A = A_1 \cup A_2 \cup \ldots \cup A_n$ where n is a positive integer, then $|R| = $?
   (c) [8%] How many relations on A are equivalence relations that satisfy $b \in [h]$?

3. [10%] Construct a state diagram for a finite state machine with the input alphabet $I = \{a, b\}$ and the output alphabet $O = \{0, 1\}$ that recognizes all strings in $\{b, a\}^*\{bb\} \cup \{a, b\}^*\{aa\}$. Note that the state diagram should contain 3 states only or get no credits.

4. Let $f(x) = x/(1-x)$ and $g(x) = 1/(1-x^2)$.
   (a) [5%] If $f(x)$ is the generating function of a sequence, what is the sequence?
   (b) [5%] If $g(x)$ is the exponential generating function of a sequence, what is the sequence?
   (c) [5%] Use $f(x)$ and $g(x)$ to express the generating function for the number of integer solutions for the equation, $c_1 + 2c_2 = 100$ with $c_1 \geq 0$ and $c_2 \geq 1$, such that the answer is the coefficient of x^{100}.

5. Solve the following recurrence relations:
   (a) [5%] $a_n = a_{n-1} + 2a_{n-2}$, $n \geq 2$, $a_0 = -3$, $a_1 = 0$
   (b) [6%] $a_n + 2a_{n-1} - 4a_{n-2} - 8a_{n-3} = 0$, $n \geq 3$, $a_0 = 0$, $a_1 = -6$, $a_2 = 8$
   (c) [6%] $a_n - 5a_{n-1} = 2(7^n)$, $n \geq 1$, $a_0 = 10$
   (d) [8%] $a_n + 2a_{n-1} + 2a_{n-2} = 0$, $n \geq 2$, $a_0 = 1$, $a_1 = 1$. Note that the solutions should contain real numbers only or get no credits.

6. [10%] Let $G = (V, E)$ be a loop-free graph with $|V| = n \geq 2$. For all $x, y \in V$ and $x \neq y$, $\deg(x) + \deg(y) \geq n-1$. Prove that G is connected.

7. [10%] Prove that an integral domain $(T, +, \cdot)$ is a field if it is finite.

1. (15%) Solve the general and particular solutions of the following ordinary differential equation (ODE).

$$y' \cos^2 x + 3y = 1, \quad y(\tfrac{1}{4}\pi) = \tfrac{4}{3}$$

2. (25%) Solve the given nonhomogeneous linear ODE only by the method of variation of parameter. Show the details of your work. [Hint : use $y = u \cdot x^{-1/2}$]

$$\left(x^2 D^2 + xD + \left(x^2 - \tfrac{1}{4} \right) I \right) y = x^{3/2} \sin x$$

3. (15%) Find the coefficient of Fourier series in the signal $x(t) = \sin(10t)\left(\sum_{k=-3}^{3} \dfrac{1}{1+jk} e^{jkt} \right).$

[Note: the coefficient of each Fourier series can be represented in the fractional form.]

4. (9%) Find the solution of the following equations only by Cramer's law.

$$8x_1 - 4x_2 + 3x_3 = 0$$
$$x_1 + 5x_2 - x_3 = -5$$
$$-2x_1 + 6x_2 + x_3 = -4$$

5. (16%) The matrix $A = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -2 & 0 \\ 0 & 0 & 3 \end{bmatrix}$,

5.1 (6%) Please find the eigenvalues.
5.2 (6%) Please find, for each eigenvalue, the corresponding eigenvector of the matrix.
5.3 (4%) Please also find the orthogonal matrix that diagonalizes the matrix.

6. (20%) Use the Laplace transform to solve the initial value problem of the following ODE. Show the details of your work

$$y_1' = 2y_1 + y_2, \quad y_2' = 4y_1 + 2y_2 + 64tu(t-1), \quad y_1(0) = 2, \quad y_2(0) = 0,$$