

Comprehensive Exam: Sample Paper

Instructions to candidates

1. Exam consists of three sections:
 - Part A: Analysis of Algorithms (6 Questions)
 - Part B: Architecture and Computer Organization (6 Questions)
 - Part C: Operating Systems (6 Questions)
2. Total Exam duration is 3 hours, 1 hour per section
3. Students will complete the exam one section at a time, with a 10 minute break in between.
4. You must write your name and your student ID on every sheet (question sheet and answer sheet).
5. Any use of cell phone or wireless device is banned during the exam. You cell phones should switched off and be put away in bags. Failing to comply with this provision will be considered using unfair means in the examination and appropriate action will be taken.
6. Calculators will be made available by the proctor
7. Bathroom break rules:
 - a. One person at a time
 - b. 10 min maximum

Part A: Algorithm Analysis and Design

Duration: 60 Minutes

Maximum Marks: 100

Name _____

Student ID: _____

Answer any 5 questions

1. **[20 points]** The boss of a company is planning a party. The company has a hierarchical structure; that is, the supervisor relation forms a tree rooted at the president. The personnel office has ranked each employee with a conviviality rating, which is a real number. In order to make the party fun for all attendees, the president does not want both an employee and his or her immediate supervisor to attend. You are given a tree that describes the structure of the company, each node of the tree holds, in addition to the pointers, the name of an employee and his or her conviviality ranking. Describe an algorithm to make up a guest list that maximizes the sum of the conviviality ratings of the guests. What is the running time?
2. **[20 points]** The diameter of a graph is the largest of all shortest-path distances in the graph. Give an algorithm to compute the diameter of a graph.
3. **[20 points]** FCC is considering the new radio station channel applications in a major metropolitan area. The stations scatter around the area such that some of them are clustered in close vicinities whereas some of them are far apart. It is well known that the stations that are close to each other should be assigned different channels as they will interfere with each other otherwise but those that are far apart from each other do not suffer from this issue. FCC wants to minimize the total number of channels assigned in this area. Prove this to be a NPC problem by reduction to a known NPC problem.
4. **[20 points]** You are given a collection of n bolts of different widths and n corresponding nuts. You are allowed to try a nut and bolt together, from which you can determine whether the nut is larger than the bolt, smaller than the bolt, or matches the bolt exactly. However, there is no way to compare two nuts together or two bolts together. The problem is to match each bolt to its nut. Design an algorithm for this problem with average running time $\theta(n \log n)$.
5. **[20 points]** Given a sequence of n real numbers $A(1) \dots A(n)$, determine a contiguous subsequence $A(i) \dots A(j)$ for which the sum of elements in the subsequence is maximized.
6. **[20 points]** A palindrome is a nonempty string over some alphabet that reads the same forward and backward. Examples of palindromes are all strings of length 1, civic, racecar, and aibohphobia (fear of palindromes).
Give an efficient algorithm to find the longest palindrome that is a subsequence of a given input string. For example, given the input character, your algorithm should return carac. What is the running time of your algorithm? 20 points

Part B: Computer Organization and Architecture

Duration: 60 Minutes

Maximum Marks: 100

Name _____

Student ID: _____

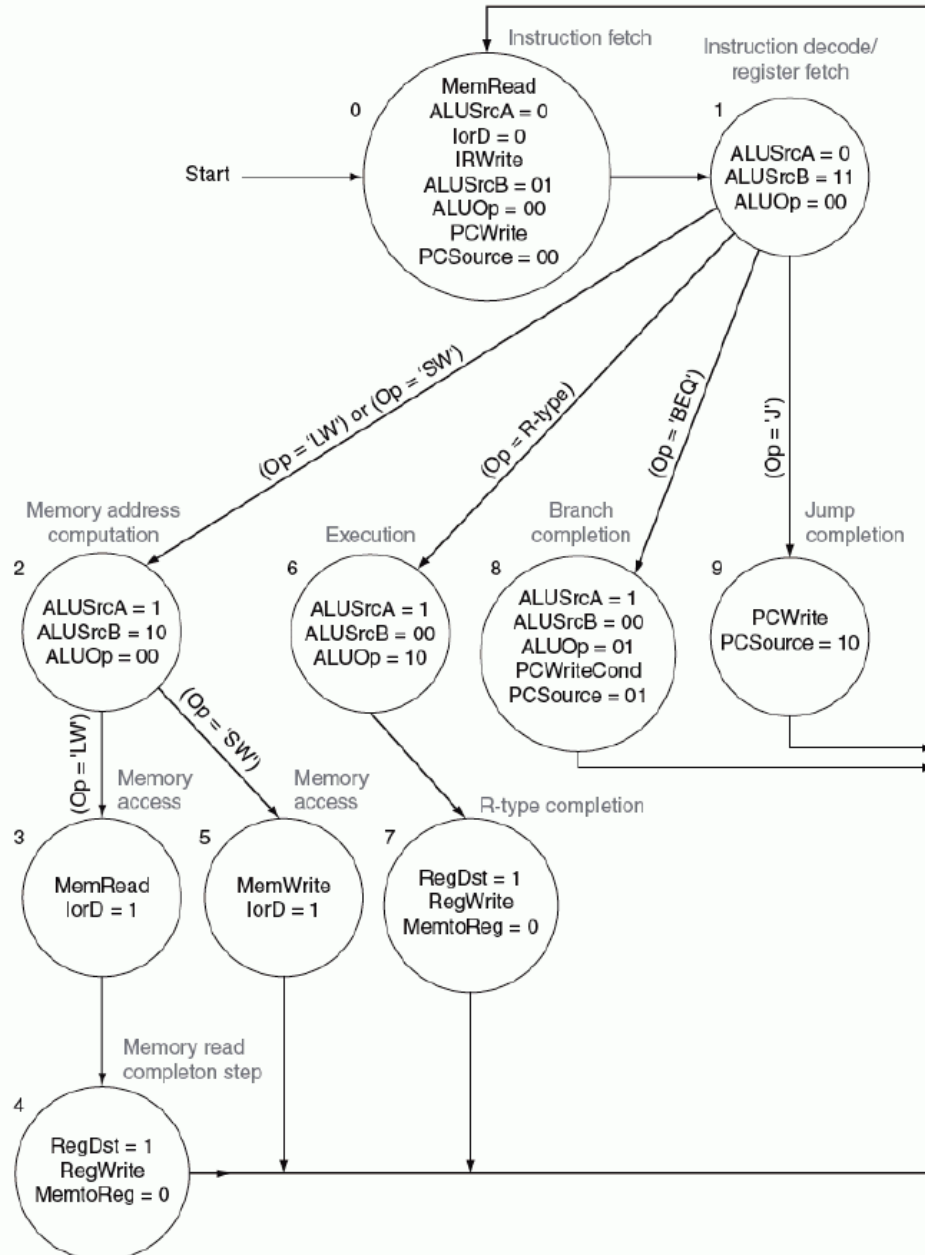
Answer any 5 questions

- [20 points]** Consider a machine that has CPI of 1.0 and a 5 GHz processor with a 5% miss rate, 100ns DRAM access, adding 2nd level cache with 5ns access time decreases miss rate to .5%, what is the CPI for the improved system? How many times is the new system more efficient than the original one?
- [20 points]** Consider program P, which runs on a 1GHz machine in 10 seconds, An optimization is made to P, replacing all instances of multiplying a value by 4 (mult X, X, 4) with two instructions that set x to x + x twice (add X, X; add X, X) Call this new optimized program P'. The CPI of a multiply instruction is 4, and that of add is 1. After recompiling, the program runs in 9 seconds on machine M. How many multiples were replaced by the new compiler?
- [20 points]** Two different compilers are being tested for a 4 GHz. machine with three different classes of instructions: Class A, Class B, and Class C, which require one, two, and three cycles per instruction (respectively). Both compilers are used to produce code for a large piece of software. The first compiler's code uses 5 million Class A instructions, 1 million Class B instructions, and 1 million Class C instructions. The second compiler's code uses 10 million Class A instructions, 1million Class B instructions, and 1 million Class C instructions. Which sequence will be faster according to MIPS? Which sequence will be faster according to execution time?
- [20 points]** How many total bits are required for a direct-mapped cache with 16 KB of data and 8-word blocks, assuming a 32-bit address?
- [20 points]** Consider the simple five-stage pipeline (Fetch, Decode, Execute, Memory, Write-back)
ADD R1, R2, R3
SUB R4, R1, R5
For the above pairs of instructions, determine how many cycles it would take to complete execution.
 - Without forwarding
 - With forwarding
- [20 points]** Two important parameters control the performance of a processor: cycle time and cycles per instruction. There is an enduring trade-off between these two parameters in the design process of microprocessors. While some designers prefer to increase the processor frequency at the expense of large CPI, other designers follow a different school of thought in which reducing the CPI comes at the expense of lower processor frequency. Consider the following machines, and compare their performance using the following instruction mix: 25% loads, 13% stores, 47% ALU instructions, and 15% branches/jumps. Assume the unmodified multi-cycle datapath and finite state machine (see figure below).
M1: The multicycle datapath is designed with a 1 GHz clock
M2: A machine like M1 except that register updates are done in the same clock cycle as a memory read of ALU operation. Thus in the finite state machine, states 6 and 7 and states 3

and 4 are combined. This machine has an 3.2 GHz clock, since the register update increases the length of the critical path.

M3: A machine like M2 except that effective address calculations are done in the same clock cycle as a memory access. Thus states 2, 3, and 4 can be combined, as can 2 and 5, as well as 6 and 7. This machine has a 2.8 GHz clock because of the long cycle created by combining address calculation and memory access.

Find out which of the machines is fastest. Are there instruction mixes that would make another machine faster, and if so, what are they?



Part C: Operating Systems

Duration: 60 Minutes

Maximum Marks: 100

Name _____

Student ID: _____

Answer any 5 questions

1. [20 points] Consider the following set of processes, with the length of the CPU burst and given in milliseconds:

Process	Burst Time	Priority
P1	4	3
P2	7	2
P3	2	6
P4	5	1
P5	4	2

The processes are assumed to have arrived in the order P1, P2, P3, P4, P5 all at time 0.


- a) Draw four Gantt chars that illustrate the execution of these processes using the following scheduling algorithms: FCFS, SJF, nonpreemptive priority (a smaller priority number implies a higher priority), and RR (quantum =2).
- b) What is the turnaround time of each process for each of the scheduling algorithms in part a?
- c) What is the waiting time of each process for each of these scheduling algorithms?
2. [20 points] Assuming that the system distinguishes between four types of resources, (A, B, C and D), there are 4 processes, P1, P2, P3 and P4. Following table illustrates the allocated, maximum and available resources:

Allocated	Maximum	Available
A B C D	A B C D	A B C D
P1 1 2 2 1	P1 4 3 2 2	3 1 1 2
P2 1 0 3 3	P2 1 2 3 4	
P3 1 1 1 0	P3 1 1 5 0	
P4 2 1 0 3	P4 3 2 1 5	

Run the Banker’s algorithm and determine if the system is in a safe state? If possible provide a safe sequence. Show all your steps.


3. [20 points] Assume a system has a TLB hit ratio of 95%. It requires 10 nanoseconds to access the TLB, and 90 nanoseconds to access main memory. What is the effective memory access time in nanoseconds for this system?
4. [20 points] Given five memory partitions of 100Kb, 500Kb, 200Kb, 300Kb, 600Kb (in order), how would the first-fit, best-fit, and worst-fit algorithms place processes of 417 Kb, 112 Kb, 212 Kb, and 426 Kb (in order)? Which algorithm makes the most efficient use of memory?
5. [20 points] Suppose we have the following page accesses: 1 2 3 5 4 2 1 5 4 2 6 3 1 2 4 7 5 4 3 5 and that there are four frames within our system. Using the FIFO replacement algorithm, LRU replacement algorithm, and optimal replacement algorithm respectively, what is the number of page faults for the given reference string for each replacement algorithm? Draw the frame replacement for each frame reference.

FIFO reference string




the number of page faults:

LRU reference string



the number of page faults:

Optimal replacement algorithm reference string



the number of page faults:

6. Explain what goes wrong in the following variation of Peterson's algorithm:

```
Process Pi:
do {
    flag[i] = TRUE;
    turn = i;
    while (flag[j] && turn == j);
    <critical section>
    flag[i] = FALSE;
    <remainder section>
} while (TRUE);
```