

8/4
Reg. No. :

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Question Paper Code : 52891

B.E./B.Tech. DEGREE EXAMINATIONS, APRIL/MAY 2019.

Seventh Semester

Electronics and Communication Engineering

EC 6009 – ADVANCED COMPUTER ARCHITECTURE

(Regulation 2013)

Time : Three hours

Maximum : 100 marks

Answer ALL questions.

PART A — (10 × 2 = 20 marks)

1. Define the terms Module Reliability and Module Availability.
2. State Amdahl's law.
3. How many bits are in the (0, 2) branch predictor with 4K entries? How many entries are in a (2, 2) predictor with the same number of bits?
4. What is meant by dynamic branch prediction?
5. Specify the benefits of data level parallelism over Instruction level parallelism.
6. Specify an architecture that supports data-level parallelism.
7. How does distributed shared memory supports thread level parallelism?
8. What is memory consistent model?
9. What are the standard levels of RAID?
10. Differentiate between throughput and response time.

PART B — (5 × 13 = 65 marks)

11. (a) (i) Explain how management of input-output device interface improve the performance of a computing processor. (8)
(ii) How to measure the performance of a desk top system using SPEC benchmark? (5)

Or

- (b) Discuss in detail about the quantitative principles of computer design. (13)

12. (a) (i) Explain the basic compiler techniques for exploiting ILP. (9)
(ii) Briefly explain static branch prediction technique. (4)

Or

- (b) How to overcome data hazards with dynamic scheduling? Explain in detail. (13)
13. (a) (i) Compare super scalar architecture with super pipeline architecture. (6)
(ii) Explain loop level parallelism. (7)

Or

- (b) Explain vector architecture in detail.
14. (a) How to measure the performance of symmetric shared multiprocessors? Explain. (13)

Or

- (b) Explain the models of memory consistency. Discuss the issues in each model with examples. (13)
15. (a) Explain cache optimization techniques. (13)

Or

- (b) Suppose a processor sends 80 disk I/Os per second, these requests are exponentially distributed, and the average service time of an older disk is 25 ms.

Answer the following questions:

- (i) On average, how utilized is the disk? (4)
(ii) What is the average time spent in the queue? (4)
(iii) What is the average response time for a disk request, including the queuing time and disk service time? (5)

PART C — (1 × 15 = 15 marks)

16. (a) Explain the support of ILP to exploit thread level parallelism with an illustration.

Or

- (b) Explain performance issues in distributed shared memory.