Keywords and Review Questions

lec1: Keywords: ISA, Moore's Law
Review Questions:
Q1. Who are the people credited for inventing transistor?
Q2. In which year IC was invented and who was the inventor?
Q3. What is ISA? Explain its role of ISA in a modern computer system?
Q4. Distinguish between architecture and organization of a computer system.
Q5. What is Moore's Law? How do you interpret it in the present day context?
Q6. Explain how instruction level parallelism can be used to improve the performance of a computer system.

lec2: Keywords: MIPS, SPEC INT
Review Questions:
Q1. How do you define performance in terms of time?
Q2. Why MIPS is considered as a meaningless information about processing speed?
Q3. How do you measure performance of a computer using benchmarks?
Q4. What is SPEC of a computer system?
Q5. What is Amdhal's Law? Using Amdahl's law, compare speedups for a program that is 98% vectorizable for a system with 16, 64, 256, and 1024 processors.
Q6. A workstation uses a 1.5 MHz processor with a claimed 10 MIPS rating to execute a given program mix. Assume a one-cycle delay for each memory access.
   (a) What is the effective CPI of the Computer?
   (b) Let the processor speed is upgraded to 30 MHz clock keeping the memory system unchanged. If 30% of the instructions require one memory access and another 5% require two memory accesses per instruction, what is the performance of the upgraded processor with a compatible instruction set and equal instruction counts in the given program mix?

lec3: Keywords: Stack architecture, RISC architecture
Review Questions:
Q1. What are the things ISA defines?
Q2. What are the programmer visible parts of a processor?
Q3. What are the ISA design choices?
Q4. Distinguish between stack based and accumulator based architecture.
Q5. What are the broad classifications of the operations of a computer?
Q6. What do you mean by instruction format? What role does it play in deciding the performance of a processor?
Q7. Distinguish between CISC and RISC architectures.
Q8. How RISC architecture leads to performance improvement?

lec4: Keywords: Register organization, Instruction format
Review Questions:
Q1. Give the register organization of the MIPS processor.
Q2. What are the data types supported by MIPS processor?
Q3. What are the addressing modes supported by MIPS processor.

Q4. Give the instruction formats of MIPS processor?

Q5. Give the data transfer instructions provided by MIPS processor.

lec5: Keywords: Data path, Control unit

Review Questions:

Q1. Give the abstract view of the data path of the MIPS processor?

Q2. What is the function of the control unit of a processor?

Q3. Give the realization of the main control unit.

Q4. Show the 4-bit ALU control input realizing different ALU operations.

Q5. Give the schematic diagram of the data path and control unit supporting R-type, memory and branch if equal instructions.

lec6: Keywords: Synchronous pipeline; Speedup

Review Questions:

Q1. Define pipelining?

Q2. What are the primary requirements to implement pipelining?

Q3. Give the basic scheme for implementing a synchronous pipeline.

Q4. Give the basic scheme for implementing an asynchronous pipeline.

Q5. How do you determine the maximum clock frequency of a pipelined processor?

Q6. Derive the expression of speedup of a k-stage pipelined processor.

Q7. Find out the throughput of a k-stage pipelined processor.

Q8. Give the implementation of a fixed point pipelined multiplier.

Q9. Give the implementation of a floating point pipelined adder.

lec7: Keywords: Pipeline register, Throughput

Review Questions:

Q1. Give the 5-stage pipelined implementation of the MIPS data path.

Q2. What operations are performed by each stage of the 5-stage pipelined processor?

Q3. How speedup is achieved in a pipelined processor? Find out the speedup factor of a 5-stage pipelined processor to execute 100 instructions, assuming that there is no hazard in executing the instructions.

Q4. Show the data path of the 5-stage RISC processor including the pipeline registers.

Q5. Why is it easy to implement pipelining for a RISC processor compared to a CISC processor?

Q6. Assume that a multiple cycle RISC implementation has a 20 ns clock cycle, loads take 5 clock cycles and account for 40% of the instructions, and all other instructions take 4 clock cycles. If pipelining the machine adds 1 ns to the clock cycle, how much speedup in instruction execution rate do we get from pipelining?

Q7. What factors limit the number of pipeline stages?

Q8. Does the performance of a pipelined processor increases linearly with the number of stages?

Q9. Consider a 7-stage pipelined processor. In the first stage, an instruction is *fetched*. In the second stage, the instruction is *decoded* and branch target address is computed for branch instructions. In the third stage, the *branch outcome* is evaluated. Assume that 25% of all branches are unconditional branches. Of the conditional branches, on an average 80% turn out to be untaken. Compute the average pipeline stall cycles per branch

instruction under pipeline stall, conditional untaken, and delayed branch schemes. Ignore the structural and data hazards. For delayed branch scheme, assume that suitable successor instructions are always found.

Lec8: Keywords: Data hazards, Bypassing
Review Questions:
Q1. What is hazard in a pipelined processor? Classify different types of hazards.
Q2. What is structural hazard? What are the popular methods to avoid structural hazard?
Q3. Explain how data dependences lead to hazards?
Q4. Why is it difficult to detect data hazards involving memory?
Q6. Explain different types of dependences with example..
Q7. Classify different types of data hazards.
Q.8. Why WAR and WAW hazards cannot occur in the 5-stage MIPS pipeline?
Q9. Consider the following Code sequence of MIPS where each instruction carries their usual meaning

$$
\begin{array}{lll}
\text{ADD} & \text{R2, R5 , R4} & \text{;Register R2} \leftarrow \text{R5} + \text{R4} \\
\text{ADD} & \text{R4, R2, R5} & \text{;R4} \leftarrow \text{R2+ R5} \\
\text{SW} & \text{R5, 100(R2)} & \text{;M [100 + (R2)]} \leftarrow \text{R5} \\
\text{ADD} & \text{R3, R2, R4} & \text{;R3} \leftarrow \text{R2 + R4}
\end{array}
$$

Enlist each of the DATA Dependencies present in this code along with its type.
Specify which of the DATA hazards present in the above code sequence can be resolved by forwarding. Justify your answer.
Q10. What is the execution time (in terms of clock cycles) of the following instruction sequence on a pipeline processor having five stages, namely Instruction Fetch, Instruction Decode, Register Read, Execute and Write Back. Assume that no bypassing is done.

ADD r3, r4, r5
SUB r7, r3, r9
MUL r8, r9, r10
ASM r4, r8, r12

Q11. Consider the following code sequence. Identify different types of hazards present and find out its execution time on a 7-stage pipeline with 2-cycle latency for non-branch instructions and 5-cycle latency for branch instructions. Assume that branch is not taken.

BNE   r4, #0, r5
DIV   r2, r1, r7
ADD   r8, r9, r10
SUB   r5, r2, r9
MUL   r10, r5, r8

Lec9: Keywords: Instruction scheduling, Loop unrolling
Review Questions:
Q1. How forwarding can be used to overcome stalls in case of data hazards?
Q2. What additional hardware are required to implement forwarding or bypassing?
Q3. Give the example of a data hazard which cannot be overcome by bypassing.

Q4. How scheduling of instructions (reorder) can be done to avoid or to reduce the number of stalls?

Q5. How loop unrolling improves instruction level parallelism to avoid stalls?

Lec10: Keywords: Loop-independent dependence, Register pressure

Review Questions:

Q1. What problems do we face by loop unrolling?

Q2. How stalls are avoided using software pipelining?

Q3. What are the basic steps to implement software pipelining?

Q4. Compare loop unrolling with software pipelining.

Q5. How can you achieve CPI less than one?

Q6. Unroll the following loop program three times and use software pipelining. Compare the execution time (including stalls as per the following table) of the unrolled program with that of the program obtained by software pipelining.

```
loop: L.D       F0, 0(R1)
      MUL.D     F4, F0, F2
      S.D       F4, 0(R1)
      DSUBUI    R1, R1, #8
      BNE       R1, R2, loop
```

Source Ins	User Ins	Latency
FP ALU Op	FP ALU Op	3
FP ALU Op	Store Double	2
Load Double	FP ALU Op	1
Load Double	Store Double	0

Lec11: Keywords: ILP, VLIW
Review Questions:
Q1. What is the maximum value of CPI possible with pipelined processors?
Q2. How higher ILP is provided in VLIW processors?
Q3. How multiple operations are performed in a VLIW processor in a single cycle?
Q4. Give an overview of the Transmeta's Crusoe processor?
Q5. Explain the function of the code morphing software.
Q6. How power dissipation is reduced in a VLIW processor?
Q7. Highlight the problems associated with VLIW processors?

Lec12: Keywords: Superscalar, Superpipelined
Review Questions:
Q1. Explain the basic principle of operation of a superscalar processor. Explain how a superscalar architecture helps to achieve CPI less than 1.
Q2. Obtain the speedup of a pipelined superscalar processor with respect to simple pipelined processor?
Q3. How higher ILP is obtained in a Superscalar processor?
Q4. Compare VLIW with superscalar processor?
Q5. What is superpipelined organization?
Q6. What are the limitations of a scalar pipelined processor.
Q7. Explain the operation of the X86 superscalar processor?
Q8. How a compiler would schedule the following code sequence for execution on a VLIW processor with 4 execution units, each of which can execute any instruction type? Load operations have 2-cycle latency, and all other operations have 1-cycle latency. Assume that the compiler examines all possible instruction orderings to find the best schedule. Include the NOPs for unused operations.

```
ADD r1, r2, r3
SUB r5, r4, r5
LD   r4, (r7)
MUL r4, r4, r4
ST   (r7), r4
LD   r9, (r10)
LD   r11, (r12)
ADD r11, r11, r12
MUL r11, r11, r11
ST   (r12), r11
```

Q9. Explain the function of the code morphing software used for Transmeta's Crusoe processor.

Lec13: Keywords: Dynamic scheduling, Out-of-order execution
Review Questions:
Q1. Explain the basic operation of the dataflow computer.
Q2. Distinguish between static with dynamic instruction scheduling used to reduce the number of stalls in a pipelined processor.
Q3. What are the advantages of dynamic scheduling of instructions?
Q4. Explain how scoreboarding is used to achieve out-of-order execution of instructions?
Q5. How WAR hazard is overcome in scoreboarding?
Q6. Give a schematic diagram of the scoreboard for MIPS processor?
Q7. Explain the operation of the four stages of scoreboard control.
Q8. What are the three parts of the scoreboard database?

Lec14: Keywords: Dynamic register renaming, Reservation stations
Review Questions:
Q1. How register naming is performed in Tomasulo's algorithm.
Q2. What are the key features of Tomasulo's algorithm?
Q3. Explain Tomasulo's scheme with the help of a schematic diagram?
Q4. How hazard detection is performed in a distributed manner in Tomasulo's algorithm?
Q5. Explain the operation of the reservation stations?
Q6. Explain the operation of the three stages of Tomasulo's algorithm.
Q7. What are the three parts of the Tomasulo's algorithm?

Lec15: Keywords: Branch penalty, Delayed branch
Review Questions:
Q1. Explain how control hazard takes place.
Q2. Why control hazard leads to greater performance loss?
Q3. How can you reduce branch penalty from 3 cycles to 1 cycle?
Q4. Explain how 'predict-not-taken' works to reduce branch penalty.
Q5. Explain how scheme is used to reduce branch penalty.
Q6. Explain how instructions are scheduled in the branch delay slot.

Lec16: Keywords: Branch prediction, correlating branch-predictor
Review Questions:
Q1. What is the basic idea behind branch prediction to reduce branch penalty?
Q2. Explain how direction-based prediction is used to reduce branch penalty.
Q3. Explain how history-based prediction is used to reduce branch penalty.
Q4. Explain the 2-bit dynamic branch prediction scheme.
Q5. Explain how correlating branch-predictor works.

Lec 17: Keywords: Tournament predictor, Branch target buffer
Review Questions:
Q1. Explain how tournament branch-predictor works.
Q2. What is the role of the branch prediction buffer to reduce branch penalty.
Q3. Explain how the return address predictor is used.
Q4. Explain the branch prediction scheme used in the Pentium processor.
Q5. Explain the branch prediction scheme used in the DEC Alpha processor.

Lec18: Keywords: Data flow architecture, Dynamic register renaming,
Review Questions:
Q1. How hardware rearranges the instruction execution to reduce stalls?
Q2. How hardware scheduling allows the processor to tolerate unprecedented delays such as cache misses?
Q3. How hardware scheduling allows code that was compiled for one pipeline in mind to run efficiently on a different pipeline?
Q4. Explain how Tomasulo's Scheme Overlap Iterations of Loops?
Q5. How hazard detection logic is distributed using reservation stations?
Q6. How Tomasulo's algorithm supports superscalar execution?

Lec19: Keywords: Hardware-based speculation, Reorder buffer
Review Questions:
Q1. How Tomasulo's algorithm supports hardware speculation?
Q2. How Tomasulo's algorithm separates the process of completion of execution from instruction commit?
Q3. Explain the function of the reorder buffer.
Q4. How exceptions are handled in using Tomasulo's algorithm?
Q5. Explain the operation of the four steps of speculative execution?
Q6. Discuss the advantages of speculation.

Lec21: Keywords: Hierarchical memory organization, Temporal locality
Review Questions:
Q1. How the gap in performance between memory and CPUs varies over time?
Q2. What are the basic objectives of hierarchical memory organization?
Q3. How the access time, cost per byte, memory size, transfer bandwidth and unit of transfer change as you go from lower level to higher level of memory?
Q4. Explain the inclusion property satisfied by memories of different levels.
Q5. Explain the coherence property satisfied by memories of different levels.
Q6. Explain different types of locality of references used in hierarchical memory organization.
Q7. Why hierarchical memory organization is crucial for temporary multi-core processors?
Q8. In a two-level memory hierarchy, if the top-level memory has an access time of 8 ns and bottom-level memory has an access time of 60 ns, find out the required hit rate of the top-level memory to get an average memory access time of 10 ns.
Q9. Assume that x1 and x2 are in the same cache block, which are in the shared state in the cache of both P1 and P2. For the sequence of events given in the following table, mention the states of each line and identify each event as a hit, true sharing miss or false sharing miss. Justify your answer. [6]

Time	P1	P2	State at P1	State at P2	Event
1	Read x1				
2		Write x2			
3	Read x1				
4		Write X2			
5		Read X2			
6	Read X2				

Lec22: Keywords: Block placement, Block identification
Review Questions:
Q1. What basic principle is used in cache memory implementation?
Q2. What are the basic objectives of hierarchical memory organization?
Q3. Explain the difference between spatial locality and temporal locality in the context of cache memory organization.
Q4. How a larger block size affects different C's of a cache memory   .
Q5. How the optimal size of cache memory is decided?
Q6. Explain the direct mapping approach used in cache memory mapping.
Q7. Explain the coherence property satisfied by memories of different levels.
Q8. Discuss the advantages and disadvantages of direct mapping used in cache memory organization.
Q9. How set-associative mapping overcomes the limitations of both direct mapping and fully associative mapping?
Q10. A cache has 64-KB capacity, 128-byte lines, and is 4-way set-associative. The system containing the cache uses 32-bit addresses.
        (i) How many lines and sets does the cache have?
        (ii) How many entries are required in the tag array?
        (iii) How many bits of tag are required in each entry in the tag array?
(iv) If the cache is write through, how many bits are required for each entry in the tag array, and how much total storage is required for the tag array if an LRU replacement policy is used replacement policy is used? What if the cache is write-back?

Q11. For the above machine, compare the number of comparators and tag bits for 8-way set-associative cache memory with that of the direct-mapped cache memory.
Q12. A direct-mapped 32-bit cache memory is assumed to have the following fields:

Tag	Index	Offset
31-12	11-05	4-0

(a) What is the cache line size?
(b) How many entries does the cache have?
(c) What is the overhead in percentage of the total cache memory size?
(d) Starting from power on following addresses are generated in sequence
        0, 4, 16, 132, 232, 160, 1024, 30, 140, 3100, 180, 2150
        (i)   How many blocks are replaced?
        (ii)   What is the hit ratio?
Q13. A set-associative cache has a block size of four 32-bit words and a set size of 2. The cache can accommodate a total of 8096 words. The main memory size is 256K×32bits. Design the cache structure and show how the processor's addresses are interpreted.

Lec23: Keywords: Tag, Replacement policy
Review Questions:
Q1. How the size of tag varies with associativity for the same size of cache memory?
Q2. Explain the replacement algorithms that can be used in set-associative mapping.
Q3. Distinguish between write through and write back approaches.
Q4. Explain the function of the update flag bit used in the write through approach.
Q5. Explain how the block size affects the miss rate for a given size of cache memory.
Q6. Distinguish between unified cache and split cache memory organization.

Q7. Give the cache memory organization of the Pentium-4 processor.

Lec24: Keywords: Alpha 21264, Average memory access time
Review Questions:
Q1. What are the three key parameters that capture memory system performance?
Q2. How the basic objectives of cache memory organization in terms of average memory access time, size and cost are satisfied in hierarchical memory organization.
Q3. What are the three key cache performance parameters?
Q4. What are the possible ways to reduce the cache hit time?
Q5. Explain the function of the valid bit of a cache memory. How the performance is improved using this bit.
Q6. How way-prediction can help to reduce the hit time of cache memory.
Q7. Explain the operation of the virtual cache.

Lec25: Keywords: Miss rate, hardware prefetching
Review Questions:
Q1. What are the basic approaches for reducing the miss rate of cache memory?
Q2. What are the three C's of cache misses?
Q3. Which parameters that affects the compulsory cache misses?
Q4. How the capacity misses varies with the cache size?
Q5. How the miss rate varies with the block size for different cache sizes?
Q6. How do you optimize the block size of a cache memory?
Q7. How miss rate is reduced by hardware prefetching of instructions?
Q8. Explain the tradeoff for the software prefetching of data.

Lec26: Keywords: Compiler optimizations, Loop unrolling
Review Questions:
Q1. Explain how you can reduce misses by merging arrays.
Q2. Explain how you can reduce misses by loop interchange.
Q3. Explain how you can reduce misses by loop fusion.
Q4. Explain how you can reduce misses by blocking.
Q5. How multilevel caches reduce miss penalty?
Q6. How miss penalty can be reduced by using write buffer?
Q7. How miss penalty can be reduced by using victim cache?
Q8. How miss penalty can be reduced by giving higher priority to read over write on miss?

Lec27: Keywords: Static RAM, Dynamic RAM
Review Questions:
Q1. Draw the 6-transistor static CMOS cell and explain its operation.
Q2. Explain the organization of a static RAM chip.
Q3. Explain the operation of a static RAM chip.
Q4. Draw a 4-transistor dynamic RAM cell and explain its operation.
Q5. Draw a 3-transistor dynamic RAM cell and explain its operation.
Q6. Draw a 2-transistor dynamic RAM cell and explain its operation.
Q7. Draw a 1-transistor dynamic RAM cell and explain its operation.
Q8. Explain the organization of a dynamic RAM chip.

Lec28: Keywords: EPROM, Flash memory
Review Questions:
Q1. Explain the basic organization of a read only memory.

Q2. What are the different types of ROM realizations possible?
Q3. Explain the operation EPROM.
Q4. Explain the operation flash memory
Q5. What are the possible ways of getting higher memory bandwidth?
Q6. How wider bus leads to reduction in miss penalty?
Q7. Explain interleaved memory organization in increasing memory bandwidth?
Q8. Explain the operation of SDRAM.

Lec29: Keywords: Virtual memory, Demand paging

Review Questions:

Q1. Why use of virtual memory is important in a computer system?

Q2. What are the key features provided by virtual memory?

Q3. Compare cache memory with virtual memory.

Q4. How an optimal page size is decided?

Q5. Explain how mapping of virtual address to physical address takes place.

Q6. Explain the address translation mechanism using page table.

Q7. What is page fault? How a page fault is serviced?

Q8. How do you select a victim page for page replacement?

Q9. What is the role of the dirty bit?

Q10. Explain the function of reference bit and protection bits of a page table entry.

Q11. A system has 48-bit virtual addresses and 128 MB of main memory. The page size is 4KB.

   (i) How many virtual and physical pages can address space support?
   (ii) How many page frames of main memory are there?

Lec30: Keywords: TLB, Demand paging

Review Questions:

Q1. Explain the function of the translation lookaside buffer (TLB).

Q2. How address translation takes place in case of TLB miss?

Q3. Draw the schematic diagram of a TLB and explain using a few sentences its function to make the address translation faster. Show different fields of the TLB.

Q4. Give an integrated diagram showing the TLB and cache operations for a logical/virtual address generated by a processor.

Q5. With the help of a flowchart explain how page faults and TLB misses are handled.
Q6. How memory management is done using user and supervisory modes?
Q7. Explain forward mapped or hierarchical page table organization.
Q8. What is an inverted page table? Explain its advantages with respect to conventional page table. How address translation takes place with the help of an inverted page table?
Q9. Explain segmentation scheme of memory mapping.
Q10. Explain the operation of different fields of the segment table.
Q11. How can you combine paging with segmentation for memory mapping?

Lec31: Keywords: Page replacement, Virtual machines

Review Questions:

Q1. Explain Pentium-II memory address translation mechanism

Q2. Discuss various page replacement policies.

Q3. Distinguish between system virtual machine and process virtual machine.
Q4. Explain the function of the Virtual machine monitor.
Q5. What are the advantages and disadvantages of virtual machines?
Q6. What must a VMM do?
Q7. What is the impact of virtual machines on virtual memory?
Q8. Explain the operation of the process virtual machine.

Lec32: Keywords: Magnetic disk, Seek time

Review Questions:

Q1. Explain the operation of magnetic disks.

Q2. Show a formatted disk and its different areas.

Q3. Distinguish between constant density and zone bit recording.

Q4. What is seek time of a disk?

Q5. What is average disk access time?

Q6. Explain the differences between the recording mechanisms in hard disk and CDROM

Q7. Why constant linear velocity is used in optical disks?

Q8. How programming and erasing is done in flash memory?

Q9. What are solid state disks?

Lec33: Keywords: RAID, Distributed parity

Review Questions:

Q1. How multiple disks work together to achieve better performance.

Q2. How RAID is used to achieve better performance?

Q3. Explain the operation of RAID level 0 used in magnetic disks.

Q4. Explain the operation of RAID level 1 used in magnetic disks.

Q5. Explain the operation of RAID level 3 used in magnetic disks.

Q6. Compare and contrast RAID level 4 with RAID level 5. Explain with an example how data is reconstructed in case of failure of a disk in an array of 9 disks using RAID-5.

Q7. Explain the operation of RAID level 6 used in magnetic disks.

Lec34: Keywords: x86 microprocessors, Multicycle decoding

Review Questions:

Q1. How pipelining is implemented in 8086 microprocessor?

Q2. What additional features are provided in 80186 and 80286 compared to 8086 processor?

Q3. What functions are performed by different stages of pipeline the Intel 486 processor?

Q4. Compare 386 with 486 in terms of cycle time and performance.

Q5. Give the block diagram of the Pentium processor and explain the function of different blocks.

Lec35: Keywords: NetBurst architecture, EPIC framework

Review Questions:

Q1. How in-order selection is performed in Pentium II/III processors?

Q2. How out-of-order execution is performed in Pentium II/III processors?

Q3. How in-order retire operation is performed in Pentium II/III processors?

Q4. Explain the NetBurst micro-architecture of the Pentium IV processor.

Q5. Explain the operation of the on-chip cache memories of the Pentium IV processor.

Q6. What is trace cache? How is it implemented in the Pentium IV processor?

Q7. What are the main ideas behind EPIC architecture? How IA-64 implements EPIC architecture?

Lec36: Keywords: Thread-level parallelism, MIMD computer architecture

Review Questions:

Q1. Distinguish between instruction-level parallelism and thread-level parallelism.

Q2. Distinguish between thread-level parallelism and process-level parallelism.

Q3. What additional hardware support is required for implementing multithreaded process?

Q4. State the benefits of multithreading.

Q5. Give a few application examples where multiple threads are inherent.
Q6. Classify the MIMD architectures.

Lec37: Keywords: Multicore architecture, Write invalidate protocol

Review Questions:
Q1. Discuss possible cache memory organizations in multicore processors.
Q2. Give the block diagram of the Intel core duo processor.
Q3. Give the block diagram of the Intel core i7 processor.
Q4. How the problem of cache coherence can be overcome by software means?
Q5. Distinguish between write invalidate and write broadcast protocols to overcome cache coherence.

Lec38: Keywords: Simultaneous multithreading, Latency hiding

Review Questions:
Q1. How is it possible to support using a single superscalar processor?
Q2. Explain coarse-grained multithreading along with its advantages and disadvantages?
Q3. Explain fine-grained multithreading along with its advantages and disadvantages?
Q4. Explain simultaneous multithreading along with its advantages and disadvantages?
Q5. What modification in the basic superscalar architecture is required to support simultaneous multithreading?

Lec39: Keywords: Symmetric multiprocessors, Cache coherence protocol

Review Questions:
Q1. Explain possible SMP organizations.
Q2. What are the advantages and limitations of SMPs?
Q3. Explain how multi-core processors provide lower power.
Q4. Explain the possible cache memory organizations in multi-core processors
Q5. Explain the cache coherence problem of shared-memory multi-core processors.
Q6. What are the possible approaches to overcome cache-coherence problem?
Q7. Explain the invalidation protocol based on snooping.
Q8. What is the meaning of each of the four states in the MESI protocol? Draw and explain the state transition diagram for a line in the cache at the initiating processor.
Q9. A four processor shared-memory system implements the MESI protocol for cache coherence. For the following sequence of memory references, show the state of the line containing the variable x in each processor's cache after each reference is resolved. All processor start out with the line containing x invalid in there cache.
Operations:
　　　Read x by processor 0
　　　Read x by processor 1
　　　Read x by processor 2
　　　Write x by processor 3
　　　Read x by processor 0

Lec40: Keywords: NUMA, Directory protocol

Review Questions:
Q1. Distinguish between UMA and NUMA model?
Q2. What are the limitations of SMPs?
Q3. What are the limitations of a snoop-based cache coherence protocol? How are they overcome in directory-based protocol?
Q4. Explain the operation of the three states of the directory protocol.

Q5. Show the state transition diagram for an individual cache block in a directory-based system.
Q6. Draw and explain the CPU-Cache state machine for directory-based protocol.
Q7. Draw the state transition diagram of the directory state machine.
Q8. Explain the Implementation issues in directory-based protocol.

Lec41: Keywords: Cluster computing, Cloud computing

Review Questions:
Q1. How computing is performed using cluster of computers?
Q2. What are the possible configurations for cluster implementation?
Q3. Explain the operation of storage area networks.
Q4. Give the cluster implementation of Google infrastructure.
Q5. What is the basic principle of operation of Grid computing?
Q6. Compare P2P computing with grid computing.
Q7. Explain the operation of cloud computing.
Q8. What are the benefits and advantages of cloud computing?