1. Cache Indexing:

Mapping	Address Space Size	Cache Size	Block Size	Tag Field	Index Field	Offset Field
Direct	1MB	8KB	4B	7	11	2
Fully Assoc.				18	--	2
2-Way				8	10	2
Direct	1MB	16KB	8B	6	11	3
Fully Assoc.				17	--	3
4-Way				8	9	3
Direct	1MB	32KB	8B	5	12	3
Fully Assoc.				17	--	3
2-Way				6	11	3
Direct	4GB	128KB	16B	15	13	4
Fully Assoc.				28	--	4
4-Way				17	11	4
Direct	4GB	512KB	16B	13	15	4
Fully Assoc.				28	--	4
4-Way				15	13	4

2. Explain the difference between a direct-mapped , set associative and fully associate cache organizations showing different fields in their memory address.

3. Consider a processor with 32 byte memory (i.e., 5 bit address) and 8 byte direct-mapped cache. Assume the cache block size is 1 byte. Table below shows the current state of the cache. For each of the memory reference in the reference sequence 10011, 00001, 00110, 01010, 01110, 11001, 00001, and 11100, identify whether the reference is a cache hit or miss. If it is a miss, identify the type of the miss (i.e., compulsory, conflict, or capacity). Answer by filling out the Table 2.

Index	V	Tag	Data
000	N		
001	Y	00	Mem(00001)
010	N		
011	Y	11	Mem(11011)
100	Y	10	Mem(10100)
101	Y	01	Mem(01101)
110	Y	00	Mem(00110)
111	N		

Table 1

Reference	H/M	Type
10011	M	Conflict
00001	H	
00110	H	
01010	M	Compulsory
01110	M	Conflict
11001	M	Conflict
00001	M	Conflict
11100	M	Conflict

Table 2

4. How is a virtual address translated to physical address by the operating system?

Page tables

If a system has 32-bit virtual address, 4-KB pages and 32-bit physical address, how much memory is consumed for keeping page tables for 100 processes?

Assume 4Byte per page table entry.

4KB pages = 12-bit for page offset. 20 bit for virtual page number.

2^{20} entries in page table

page table size = 2^{20} * 4 Byte = 4MB per process

page table size for 100 process = 100 * 4MB = 400MB

How can the memory space overhead be reduced?

Multi-level page tables

How can memory access overhead be reduced?

TLB

5. What is the maximum memory capacity supported by the following server: 2 processor sockets, each socket has 4 memory channels, each channel supports 2 dual-ranked DIMMs, and x4 4Gb DRAM chips? Assume 64-bit data bus

2 sockets * 4 memory channels * 2 * 2 rank * 16 * 4GB = 256GB

Note: “2 dual-ranked DIMMs” = total of 4 ranks. (Each DIMM has 2 ranks, front and back)

Note: with 64-bit data bus, we need 16 x4 DRAM chips. A x4 DRAM chip only supplies 4-bits.

6. PAGE TABLES. Consider the following parameters.
   Page size: 16KBytes
   Page table: four-level page table.
   The virtual page number is split in 4 fields of 9 bits each.
   Entries in all tables are 32 bits (4 bytes).

   a. Which bits of the virtual address are used to index the first level table (top level in the hierarchy):
   Answer: bits V49 to V41

   b. Which bits of the virtual address are used to index the page tables at the bottom of the hierarchy:
   Answer: bits V22 to V14

   c. What is the size of each table at all levels (in bytes)?
   2048

   d. What is the total amount of virtual memory covered by one entry of page tables at each level:
   level1: 2TB (2^{41})
   level2: 4GB (2^{32})
   level3: 8MB (2^{23})
   level4: 16KB(2^{14})

7. TLB. Continuing from the previous problem, consider the following parameters.
   TLB size: 128 entries
   TLB organization: 4-way set-associative

   a. Which bits of the virtual address are used to index the TLB:
   Answer: bits V18 to V14

   b. Which bits of the virtual address are used as tags in the TLB?
   Answer: bits V49 to V19

   c. TLB tag size:
   31 bits

Other conceptual questions can include:
Thread-level parallelism
Data-level parallelism
Main memory
Cache
GPGPUs
Datacenters