MS108: Computer System 1

Spring 2015

Homework #4

Due: Two Weeks from Assignment

TA: Ran Ye

Email: 叶冉 [happyinglife@sjtu.edu.cn]

Collaboration Policy

These homework sets will be extremely valuable as tools for learning the material and for doing well on the midterm and final. You are required to obey the following rules:
(a) Each student should write out their solution independently and in their own words.
(b) Same applies to programming assignments – you should do your own coding.
Above all, make sure that you understand the solution to these homework problems. They really are assigned to help you understand the material and be prepared for the types of problems on the midterm and final!

Q1. Memory

Consider a memory subsystem of hierarchical architecture. In this memory hierarchy, there are L1 cache, main memory (DRAM), and a hard disk (There is NO L2 cache). The L1 cache is a physical cache using physical addresses for indexing and tagging. The hard disk serves as the secondary storage and swapping device for the virtual memory. The whole subsystem applies following strategies.

L1 cache: L1 cache is on-chip unified cache. It uses write-through policy with write-allocate. The write-through from L1 cache goes to the main memory directly. The hit rate is H1 and the hit time is 1 cycle, meaning that there is no stall when the cache is hit. The percentage of read is r% and the write is w%.

Main memory and TLB: Main memory applies the virtual memory. It has a single level page table with on-chip TLB to speed up the page table lookup. Suppose the hit rate of TLB is H2 and the page hit rate is H3 (i.e., page fault rate is 1-H3). The hit time of TLB is one cycle. If we have to access the main memory for reading or writing a block or a page table entry, we need M stall cycles to complete reading/writing the physical memory. Assume that we have designed the physical memory to allow a form of interleaving so that reading a block and writing a word can be executed “in parallel” and complete in M stall cycles.

Disk: If we have to go to the disk to swap the page from the disk to the physical memory or from the physical memory to the disk, we need D cycles including the cycles needed

for memory and disk operations. The percentage of clean pages is c% and that of dirty pages is d%. Assume that, with single level page table, all the page tables are in the main memory. The page tables are in the fixed positions and will not be paged out to the disk.

Fig. 1 gives certain parts of the Memory Access Tree.

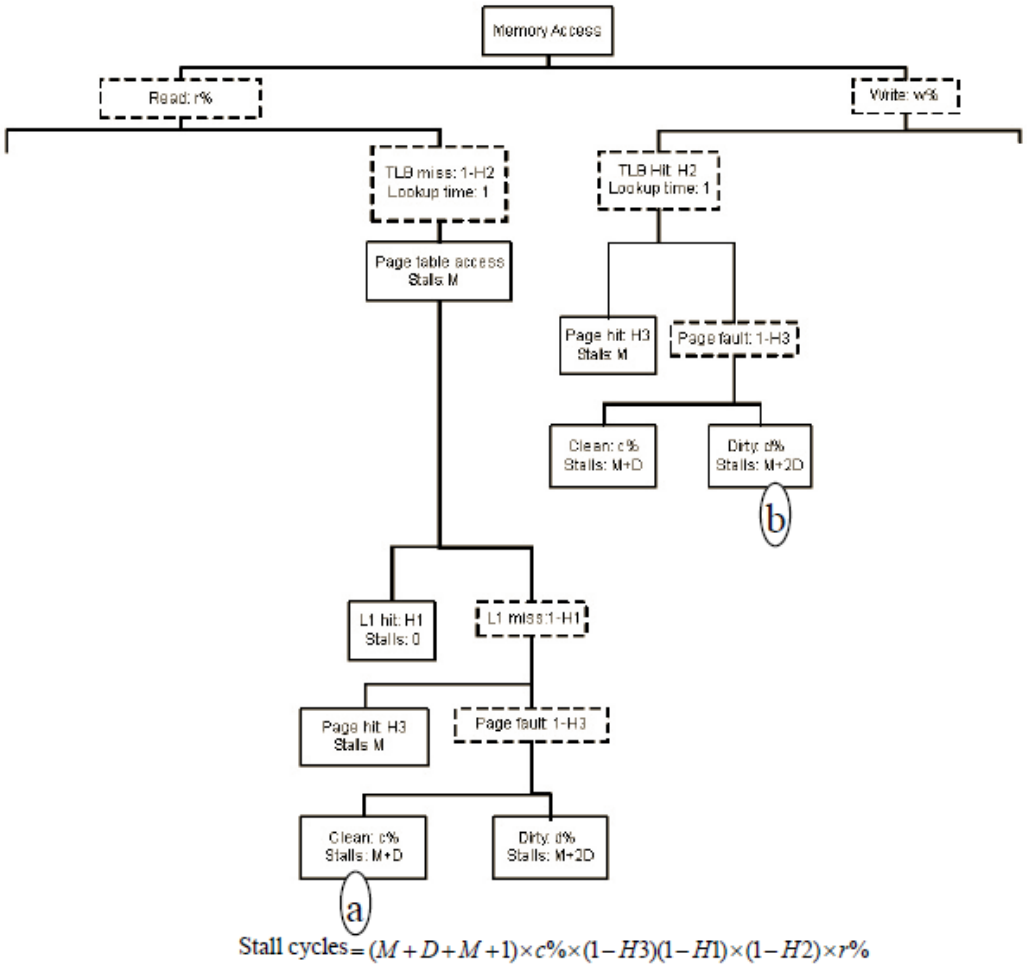

$$\text{Stall cycles} = (M+D+M+1)\times c\%\times(1-H3)(1-H1)\times(1-H2)\times r\%$$

Fig. 1: The memory access tree with the subtrees for the scenarios shown: memory read with TLB miss, memory write with a TLB hit

i. The expression of the stall cycles per memory access for case ⓐ is shown in the Fig. 1. Based on the memory access tree, give the expression of the stall cycles per memory access for the case ⓑ , the scenario of a memory write with TLB hit but page fault on a dirty page.

ii. Complete the Memory Access Tree. Either complete Fig. 1 or draw on this page.

Q2. TLB

Consider a memory system with the following parameters:

Translation Lookaside Buffer

- A total of 512 entries, organized as 4-way set associative
- LRU Replacement

L1 Cache

- 128KB, 4-way set associative, 32byte lines
- LRU replacement

Memory

- 4GB Virtual Address Space
- 512MB of physical memory
- 4KB Page Size

The following page shows a labeled diagram of the cache and TLB. Please fill in the appropriate information in the boxes below.

L1 Cache		TLB	
A=	bits	G=	bits
B=	bits	H=	bits
C=	bits	I=	bits
D=	bits	J=	bits
E=	bits		

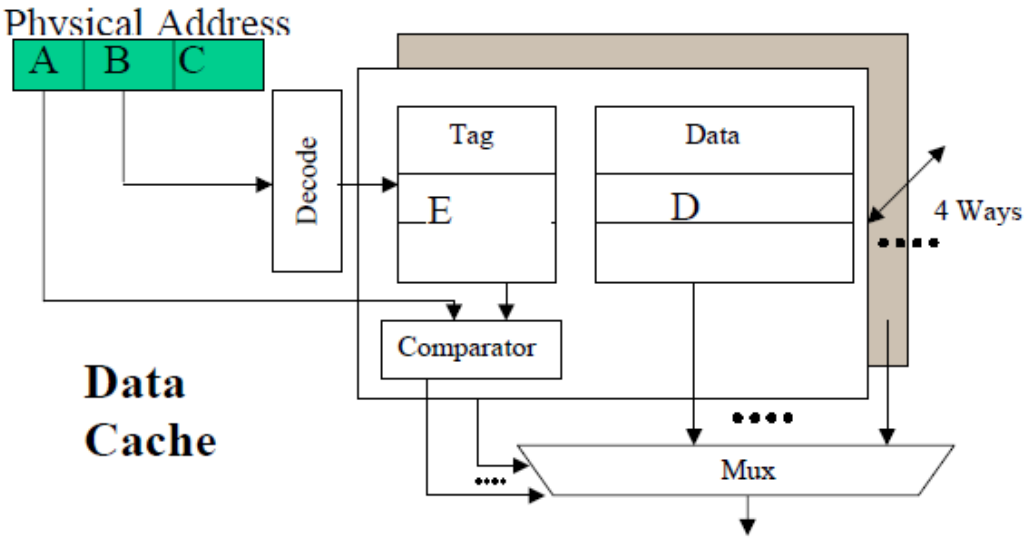

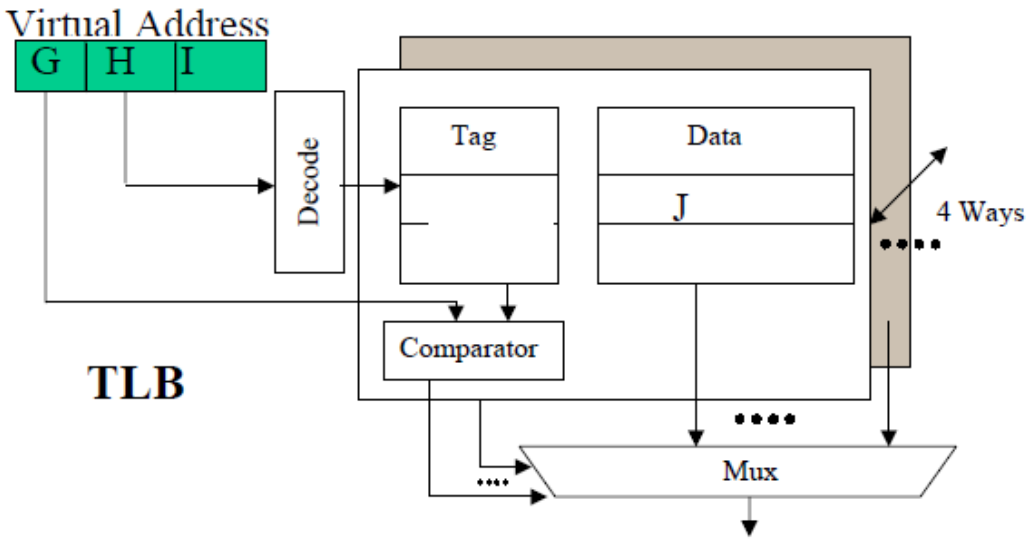

Q3. Hard Disk

Consider a group of 8 hard drives forming a RAID-5 system. Each drive has one single-sided platter and 20000 tracks on the surface. Each track has 5000 sectors, and each sector contains 512 bytes. The drive's rotational speed is 9600 rpm. The configuration of the RAID-5 system is illustrated in Figure 2.

Disk 0	Disk 1	Disk 2	Disk 3	Disk 4	Disk 5	Disk 6	Disk 7
0	1	2	3	4	5	6	P0
7	8	9	10	11	12	P1	13
14	15	16	17	18	P2	19	20
21	22	23	24	P3	25	26	27
28	29	30	P4	31	32	33	34
35	36	P5	37	38	39	40	41
42	P6	43	44	45	46	47	48
P7	49	50	51	52	53	54	55
56	57	58	59	60	61	62	P8
63	64	65	66	67	68	P9	69

Figure 2: One example of the stripe configuration of RAID 5 using 8 disks as one group.

i. How much useful storage capacity, excluding redundant information, can be provided by this set of 8 disks?

ii. Suppose disk 2 fails at certain time, and a new replacement disk is installed. Which disk/disks will be accessed to reconstruct the missing data? Specify the disk numbers and the operations (read or write) on them.

iii. When updating the information in one disk block, the RAID-5 system applies small writes to update the data block as well as the parity block. Suppose at any moment the disk array can serve TWO write requests at most. At certain time instance, there are several write requests on block 0, 6, 12, 23, 27, 28, 66, 50, and 58. Block 0 is currently being written. The other requests are waiting to be scheduled, and they are not ordered. The RAID-5 system can schedule any one to be serviced in parallel with the write of Block 0, if the request does not conflict with the write to Block 0. Among these requests, which **ones** can potentially be serviced in parallel with the write on block 0? Explain the reasons why these ones can potentially be serviced in parallel and why others cannot.
(Hint: We do not need to consider the correlation or dependence among the waiting requests.)