
EE108B

Review Session 6

Drew Hall
Stanford University
<http://eeclass.stanford.edu/ee108b>

Announcements

- Lab4 is due December 6th
- HW5 is due December 6th
- Class moved from Thursday, December 8th to Wednesday, December 7th
- Review session on December 8th during normal class time
- Final Exam is on December 14th at 12:15PM – 3:15 PM in Gates B3

Homework Hints

- Problem 1:
 - Hints not necessary
- Problem 2:
 - Identify the speed of each device involved, essentially the operations will be limited by one or two of these devices
- Problem 3:
 - Build off your answer in #2
 - Don't forget the part about the block size being a power of 2, this is important!

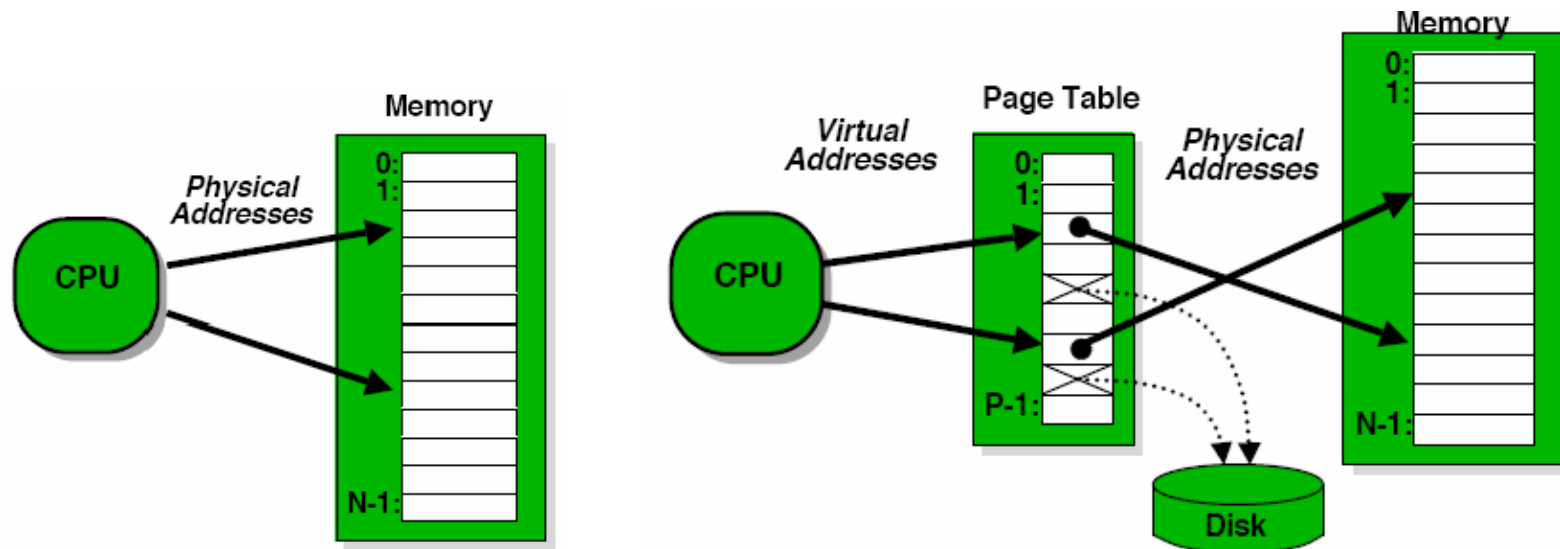
Homework Hints 2

- Problem 4:
 - A) $\text{Size} = \text{Platters} \times (\text{Heads/Platter}) \times (\text{Tracks/Head}) \times (\text{Sectors/Track}) \times (\text{Bytes/Sector})$
 - B) Break the problem up into smaller parts, calculate the number of bytes per revolution and then the number of seconds per revolution. Combine these and get the bandwidth
 - C) From the lecture, we know the formula for access time is:
 $\text{Access Time} = \text{Seek Time} + \text{Rotational Delay} + \text{Controller Overhead}$
 - E) Assume 100% utilization of the drives and transfer data in 64KB blocks
 - F) Ignore the fact that a real RAID5 configuration would perform the same operation on both drives (striped).. (it is possible...)

Operating System and Virtual Memory

- Driving Factors
 - Exploit magnetic storage medium to give the programmer the illusion of a large amount of memory regardless of the actual amount of memory
 - OS protection and shared memory
 - Memory mapped IO
 - Simplify memory management

System Comparison




Hardware to Support VM

- Extra hardware needed to support VM
 - Translation Look-aside Buffer (TLB)
 - Translate a virtual address into a physical address
 - Page Table
 - Table storing all information about the pages
 - Can be subdivided into a 2 or multilevel hierarchy

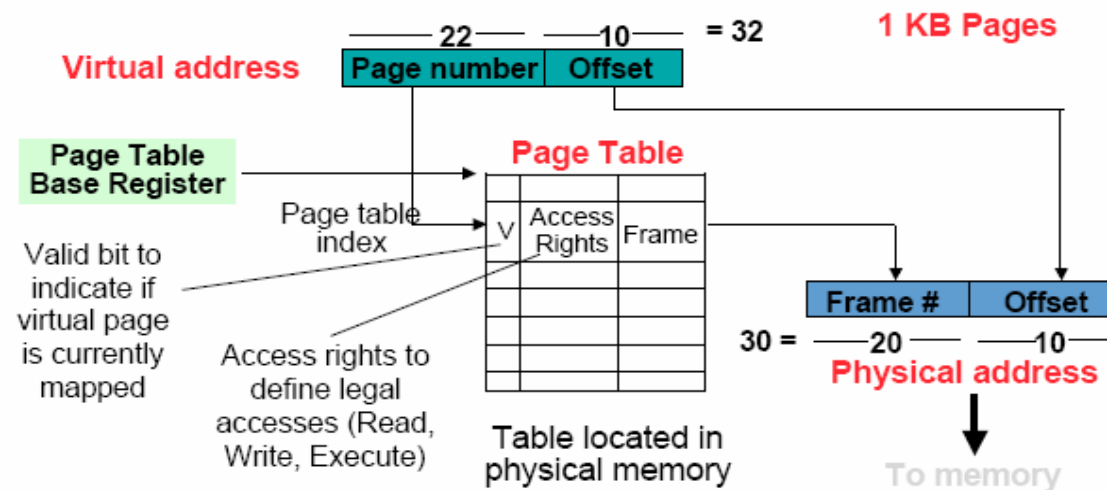
Address Translation

- Need to take a virtual address and map this address into a page and an offset
- Need to know the page size (1KB, 64 KB, 1 MB,)
- Issues to consider:
 - Do you want to store virtual or physical address in cache?
There are advantages and disadvantages to each
- Simple translation scheme:



Page Table

- Page table needs to store information about pages
 - Page Size
 - Access rights (Read, Write, Execute)
 - Valid bit (is the page currently in memory?)
 - Frame (significant bits of physical address)



Page Faults

- Occur when a page is requested that is not located in memory, similar to a cache miss
- Handled by the OS
- Expensive since page must be read from disk
- Placement strategy important for new page, similar to cache line replacement problems
- A process will typically lose control of the CPU on a page fault so that the processor can continue with other processes

Processes

- Each process has its own address space
- Processes address spaces are isolated from one another
- Information about processes stored in Process Control Block (PCB)
 - Register values
 - PC value
 - Other contextual information needed to be preserved across a context switch


Designer Considerations

- Page size
 - Affects the size of the page table, affects memory fragmentation
 - Support multiple page sizes or a single fixed size page
- Page Sharing
 - Copy On Write
- Page Access Restrictions
 - Read Only, Execute, Read / Write, ...
- TLB Size
- Memory Accesses
 - Unaligned vs Aligned
- Page Table Hierarchy

IO Devices

- Throughput
 - Average amount of time required to move data
 - Taken over a window of time
 - Amortizes the time spent waiting in the queue for the device
 - Highest throughput when the device is fully utilized
- Response Time
 - Amount of time taken to respond to a particular request
 - Fully accounts for the time spent waiting in the queue
 - What the end user typically sees
 - Fastest response time when the device is not utilized

Hard Drives



- Divided up into components
 - Sectors
 - Tracks
 - Platters
 - Cylinders

Size = Platters x (Heads/Platter) x (Tracks/Head) x (Sectors/Track)
x (Bytes/Sector)

Access Time

- Hard drive access time due to three different components
 - Seek Time – Time to move the arm to the right track
 - Rotational Latency – Time for the sector to rotate under the head
 - Transfer Time – Time required to read/write the disk
- These numbers are given in averages because absolute numbers is not possible to calculate
- Does not account for any caching that the disk may do, controller overhead, bus transfer times, etc...

Buses

- Shared communication link between devices
- Design considerations dependent on purpose of the bus
 - Point to Point vs Shared
 - Synchronous vs Asynchronous
 - Arbitration scheme
 - Split control and data lines
 - Width
 - Bandwidth

Polling vs Interrupts

- Polling
 - Processor continually checks to see if an operation has been completed
 - Fastest response time because when the event happens, the processor is ready to handle it
 - Wasted time
- Interrupts
 - Processor performs other pending tasks until it receives an interrupt
 - Slower response time due to interrupt and context switching overhead
 - Better time utilization

DMA

- Direct Memory Access (DMA) offloads memory transfers to a separate controller
- DMA Transfer Steps:
 - Processor initiates request, specifying the source, destination, and the number of bytes to transfer
 - DMA arbitrates the bus and performs the block copy operation
 - DMA controller notifies the processor when the operation is completed (via interrupt mechanism)

DMA Problems

- Address Issues:
 - A large problem with DMA transfers comes from paging. If the DMA controller works with physical addresses, then it cannot copy across pages because pages may not be consecutive in memory.
 - If the DMA works with virtual addresses, then it has to use the TLB to translate from a virtual address to a physical address. Much more complicated design
- Cache Coherency
 - Data in the cache may be more recent than the data in memory for a write back cache. Thus, the DMA controller needs to also check the cache.