Page Tables, Caches and TLBs

Page Tables

Suppose that we have a two-level page translation scheme with 4K-byte pages and 4-byte page table entries. What is the format of a 32-bit virtual address? Roughly sketch out the format of a complete page table.

Suppose we have a memory system with 32-bit virtual addresses and 4 KB pages. If the page table is full, show that a 20-level page table consumes approximately twice the space of a single level page table.

Show that the above is not necessarily true for a sparse page table.

Caching

What are the two principles of locality that make implementing a caching system worthwhile?

List and (briefly) describe three sources of cache misses.

True or False: When looking up an entry in the cache, it is not sufficient to find a matching index.

What are the advantages / disadvantages of using a Fully Associative Cache?

Briefly explain the difference between a write-through and a write-back cache.

Translation Lookaside Buffers (TLBs)

True or False: TLBs are organized as a directly-mapped cache to maximize efficiency.
True or False: The TLB in Nachos needs to always be invalidated on every context-switch.

What are the disadvantages of using the higher order bits to index into a TLB?

What are the disadvantages of using the lower order bits to index into a TLB?

What is the effective access time for TLB with 80% hit rate, 20ns TLB access time and 100 ns Memory access time (assume two-level page table that is not in L2 cache)?

Windows ReadyBoost

Windows ReadyBoost is a technology that aims to improve performance by utilizing external storage (primarily USB storage keys) as an "extended disk cache." The idea is that since it may be faster to read data from a flash drive as opposed to a hard drive, one can notice a considerable performance gain by caching some data in this USB drive. Note that USB drives have extremely fast seek times (~ 1 ms) compared to Hard Drives (~ 8 ms). However, Hard Drives have a *significantly* higher bandwidth rate.

Why do you think there is a significant difference in seek time between the two mediums?

Based on the performance information described above, what kind of I/O requests would be best served by the ReadyBoost cache? Why?

If one were to represent cache-levels via a hierarchical diagram, does ReadyBoost lie ahead of the disk cache in memory or behind the disk cache in memory?

Should ReadyBoost be a write-through cache, or a write-back cache?

Can you think of another creative way to generate an "extended disk cache?" (Hint: Fiber Optic Cables)