

The final will be 10-11:50 AM on Thursday, December 17, in ITT 328. The test will be closed-book and closed-note, except for three sheets of paper with notes. About 80% of the Final will focus on the material since the last test. About 20% will focus on tests 1 and 2 material, but will be general conceptual questions and not detailed problems.

Chapter 6. External memory

Magnetic disk- organization/format; characteristics; access time (seek time, rotational delay/latency)

RAID - levels and their relationship to error-correcting, I/O request rate, and data-transfer rate

Chapter 7. I/O

General idea of interrupts

I/O module operation and function

I/O instructions: memory-mapped I/O and isolated I/O

I/O techniques: programmed I/O, interrupt-driven I/O, and direct-memory access (DMA)

Chapter 8. OS Support and Virtual Memory

OS objectives and functions

Effects of multiprogramming on system utilization

Process control block (PCB) components

Scheduling of jobs: types of queues and their purpose

Scheduling techniques, priorities, and fairness

Memory management: swapping, fixed partition and variable partitioning

Paging: virtual memory, demand paging, page table and virtual to physical address translation, time and memory efficiency considerations, TLB (translation lookaside buffer)

Handling of large page tables: two-level page table, inverted page table

Page-replacement algorithms and their implementation (R-bits & history bits)

Frame-allocation algorithms: page-fault frequency algorithm

Segmentation: advantages and disadvantages

Segmentation and paging combined

Chapter 17: Parallel Processing (skip sections 17.5 on Clusters and 17.7 on Vector Comp.)

Motivation for multiprocessors and clusters

Flynn's Classification: SISD, SIMD, MISD, MIMD

Multiprocessor basic organizations: Communication models (message passing vs. shared memory (UMA, NUMA) and physical connection (network vs. bus)

Symmetric Multiprocessors (SMP): organization - shared memory and I/O, single bus, symmetric processors

Cache coherency, MESI protocol, (talked about with cache) process synchronization, spin locks

General idea of cache coherency in NUMAs (directory-base protocols)

Multithreading: Approaches to multithreading: fine-grain/interleaved, course-grain/blocked, simultaneous multithreading (SMT), chip multiprocessing/multicore (Figure 17.8)

Motivation of multithreading on a chip: hide stalls due to dependency, cache misses, etc. and thus increase utilization of functional units

Chapter 18: Multicore processors

Motivation for multicore processors: Pollack's rule

Multicore organizations: General structure of Intel Core Duo and Intel Core i7