Exam Review

Monitors

Deadlocks

Memory Management

Monitors

- What is a monitor?
- What is a condition variable?
- Hoare vs Mesa style monitors.
- Dining philosophers, readers and writers

Deadlocks

- What is a deadlock? Difference from starvation
- Necessary conditions for a deadlock
- Deadlock detection, avoidance, prevention
- Resource allocation graph - deadlock detection
- Concept of a safe state
- Banker's algorithm
  - Problem solving with banker's algorithm

Memory Management

Topics you should understand:

1. What is virtual memory and why do we use it?
2. Memory allocation strategies:
   - Relocation and Contiguous allocation (first-fit and best-fit algorithms)
   - Paging
   - Segmentation
   - Paged segmentation
   - Demand paging

Memory Management (cont.)

For each strategy, understand these concepts:

- Address translation
- Hardware support required
- Coping with fragmentation
- Ability to grow processes
- Ability to share memory with other processes
- Ability to move processes
- Memory protection
- What needs to happen on a context switch to support memory management

Memory Management (cont.)

Things you should be able to do:

- Given a request for memory, determine how the request can be satisfied using contiguous allocation.

- Given a virtual address and the necessary tables, determine the corresponding physical address.

Memory management strategies

- Contiguous allocation
  - static versus dynamic allocation
  - Base and limit registers
  - Best, first, worst-fit strategies

Paging and Segmentation

Topics you should understand:

- What is paging, a page, a page frame?
- What does the OS store in the page table?
- What is a TLB? How is one used?
  - Effective memory access times using a TLB
- What is segmentation? where are segment tables stored?
- How to combine segmentation and paging?
- What is demand paging?
- What is a page fault, how does the OS know it needs to take one, and what does the OS do when a page fault occurs?

Demand Paging

- What is demand paging?
- What is a page fault?

Paging (cont.)

- Page replacement algorithms. For each understand how they work, advantages, disadvantages, and hardware requirements.
  1. FIFO
  2. MIN
  3. LRU
- How do global and per-process (aka local) allocation differ?
- What is temporal locality? What is spatial locality? What effect do these have on the performance of paging?
- What is a working set?

Paging (cont.)

- What is thrashing and what are strategies to eliminate it?
- What considerations influence the page size that should be used?

Things you should be able to do:

1. Given a page reference string and a fixed number of page frames, determine how the different replacement algorithms would handle the page faults.