

Virtual Memory



Review Questions

Section 10.1

- 10.1 True or False? A program does not need to be stored in memory in its entirety.
- 10.2 True or False? A physical address space is at least as large as a virtual address space.

Section 10.2

- 10.3 When does a page fault occur?
- 10.4 True or False? In a pure demand paged system a page is never brought into memory until it is needed.

Section 10.3

- 10.5 What system call initiates copy on write?
- 10.6 True or False? The `vfork()` system call does not use copy on write.

Section 10.4

- 10.7 What is the simplest page replacement algorithm?
- 10.8 What is the name of the page replacement algorithm that operates by replacing the page that will not be used for the longest period of time?
- 10.9 What page replacement algorithm could be implemented using a stack or counters?
- 10.10 True or False? Approximation algorithms are almost always used when implementing LRU.

Section 10.5

- 10.11 What is the fundamental difference between global and local page replacement?

Section 10.6

- 10.12 What term is used to describe the situation where a process spends more time paging than executing?
- 10.13 What term is used to describe the set of pages a process is currently referencing?
- 10.14 True or False? With pure demand paging, the page fault rate is initially very high.

Section 13.5

- 10.15 True or False? Shared memory is typically not implemented using memory mapping.

Section 10.8

- 10.16 Using the buddy system, if a request for 200 KB of kernel memory is made, how much is actually allocated?
- 10.17 What is one benefit of using slab allocation.

Section 10.9

- 10.18 What is the TLB reach of a system with 4 KB page sizes and 32 entries in the TLB?
- 10.19 True or False? 4 KB is a typical page size.
- 10.20 True or False? Some systems support page sizes up to 4 MB.

Section 10.10

- 10.21 What page replacement algorithm is used by Windows?
- 10.22 Solaris uses the clock algorithm variation of LRU. How many hands does this algorithm employ?