

**Operating Systems Design**  
**Exam 2 Review: Fall 2010**  
 Paul Krzyzanowski  
 pxk@cs.rutgers.edu

1

1. Why could adding more memory to a computer make it "run faster"?

If processes don't have their working sets in memory, that leads to thrashing, which has a big impact on performance. We want system memory to be at least as large as the  $\sum$  of the working sets of all processes.

Not: *more programs can be loaded into memory (good answer for swapping systems)*

Not: *we can have a bigger buffer cache*

2

2. Some processors do not support an address space identifier in their TLB. Under what circumstances does its absence affect system performance?

Under context switching.

Without an address space identifier (ASID), context switching will require flushing the TLB, which will lead to a very low initial hit ratio.

Not: *when we have multiprogramming systems.*

Not: *a process may accidentally access another process' memory*

3

3. A hard link is having two or more file names refer to the same file. Why are hard links easy to implement on an inode-based file system (such as FFS) but difficult on a FAT-based file system (such as Microsoft's FAT32)?

- Under FFS, the directory entry points to the inode (metadata). Multiple directory entries may point to the same inode as long as the inode keeps reference (link) counts.
- Under FAT, the directory entry contains the metadata (there's no place for it in the File Allocation Table). Multiple directory entries could, in theory, point to the same starting block # but the metadata won't stay in sync (you'd also need to figure out how to do reference counting).

4

4. An inode based file system has blocks of size  $b$  bytes. Each block pointer is  $w$  bytes long. An inode entry contains  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. What is the maximum possible file size that can be supported? Be sure that your answer expresses the size in bytes and not blocks.

Lots of students got this wrong!!

5

4. An inode based file system has blocks of size  $b$  bytes. Each block pointer is  $w$  bytes long. An inode entry contains  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. What is the maximum possible file size that can be supported? Be sure that your answer expresses the size in bytes and not blocks.

Direct blocks

Capacity =  $ib$

6

4. An inode based file system has blocks of size  $b$  bytes. Each block pointer is  $w$  bytes long. An inode entry contains  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. What is the maximum possible file size that can be supported? Be sure that your answer expresses the size in bytes and not blocks.

The diagram shows an inode with a metadata section and three groups of pointers:  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. Each direct block points to a single data block. Each indirect block points to a disk block containing  $b/w$  entries, each pointing to a data block. Each double indirect block points to a disk block containing  $b/w$  entries, each pointing to another disk block containing  $b/w$  entries, each pointing to a data block.

Capacity of each indirect block:  $(b/w)b$   
Capacity:  $j(b/w)b$

4. An inode based file system has blocks of size  $b$  bytes. Each block pointer is  $w$  bytes long. An inode entry contains  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. What is the maximum possible file size that can be supported? Be sure that your answer expresses the size in bytes and not blocks.

The diagram is similar to slide 7 but includes capacity calculations for double indirect blocks. Each double indirect block points to a disk block containing  $b/w$  entries, each pointing to another disk block containing  $b/w$  entries, each pointing to a data block.

Capacity of each indirect block:  $(b/w)b$   
Capacity of each double indirect block:  $(b/w)(b/w)b$   
Capacity:  $k(b/w)^2b$

4. An inode based file system has blocks of size  $b$  bytes. Each block pointer is  $w$  bytes long. An inode entry contains  $i$  direct blocks,  $j$  indirect blocks, and  $k$  double indirect blocks. What is the maximum possible file size that can be supported? Be sure that your answer expresses the size in bytes and not blocks.

Maximum file size:  
$$ib + j(b/w)b + k(b/w)^2b = b(i + jb/w + kb^2/w^2)$$

5. How does an extent differ from a cluster?

- Cluster = fixed-size allocation unit on the disk; a logical block.
- Extent = variable-size allocation unit: multiple clusters (blocks)

Questions 6-8

6. Base and limit addressing is most useful in:

- an inverted page table.
- a single partition monoprogramming system.
- a direct mapping paging system.
- a segmentation system.

7. A hit ratio in a paging system is:

- The fraction of memory translations that are made by a memory-based page table lookup.
- The fraction of memory translations that are made by the TLB.
- The ratio of successful lookups to page faults.
- The ratio of memory writes to memory reads.

8. A Translation Lookaside Buffer (TLB) caches:

- the contents of frequently accessed memory locations.
- frequently accessed pages.
- partial page tables.
- frequently accessed page table entries.

Questions 9-11

9. A certain MMU converts 16-bit virtual address with 512-byte pages into 20-bit physical addresses. How big is the page table?

- 128 entries ( $2^7$ )
- 2048 entries ( $2^{11}$ )
- 65536 entries ( $2^{16}$ )
- 1048576 entries ( $2^{20}$ )

10. Which memory allocation strategy picks the largest hole?

- First fit
- Best fit
- Worst fit
- Next fit

11. The following page table has the same number of entries as there are page frames:

- multilevel page table
- inverted page table
- direct mapped page table
- associatively mapped page table

**Questions 12-14**

- Which page replacement algorithm best approximates LRU (least recently used)?
  - First In, First Out
  - Not Frequently Used [at each timer interrupt, increment counters on used pages – no history]
  - Clock (second chance) [pick a page that has not been referenced]
  - N<sup>o</sup> chance** [increment a per-page counter if not referenced; pick page with that hasn't been referenced over N faults]
- Suppose that you have a 32-bit address with a two-level paging hierarchy and a 4 KB page size. The top-level index table has 1024 entries. How many entries does each partial page table have?
 

32 bits		
1 <sup>st</sup> level	2 <sup>nd</sup> level	offset
10 bits: 1024 bytes (2 <sup>10</sup> )	32 - 10 - 12 = 10 bits	12 bits: 4096 bytes (2 <sup>12</sup> )

  - 1024 (2<sup>10</sup>)
  - 262144 (2<sup>18</sup>)
  - 1048576 (2<sup>20</sup>)
  - 4194304 (2<sup>22</sup>)
- Page fault frequency monitoring is a technique for:
  - managing the resident set for all processes in the system.**
  - evaluating the effectiveness of the TLB.
  - evaluating the effectiveness of the MMU.
  - optimizing the interrupt response time for page faults.

**Questions 15-17**

- An example of a block device is a:
  - Bluetooth headset.
  - Scanner.
  - Video frame buffer.
  - CD-ROM.**
- A thread should not block if it is executing in:
  - User context.
  - Interrupt context.**
  - Kernel context.
  - (b) or (c).
- On POSIX systems, a device is identified within the file system by:
  - A file name that has a unique meaning to the kernel.
  - A file name that is a hard link to the device driver.
  - A file name that points to the device table instead of to an inode.
  - An inode entry that contains a number that is the index into the device table as well as an additional parameter.**

**Questions 18-19**

- The goal of the top half of a device driver is to:
  - Delegate as much work to the bottom half as possible.**
  - Forward user requests from a system call to the device controller.
  - Initialize the device driver and register it with the kernel.
  - Keep track of the number of references to this specific driver.
- Assume that we're using logical block addresses instead of tracks and sectors to refer to disk locations and assume that low block numbers are on the outside of the disk and high block numbers are on the inside. The last block written is 10,000. The block written before that was 8,000. Our queue for read/write requests contains:
 

6000, 11000, 30000, 9000, 20000

 What is the sequence in which the requests will be scheduled with a LOOK algorithm?
  - 9000, 11000, 6000, 20000, 30000 [Shortest Seek Time First]
  - 11000, 20000, 30000, 9000, 6000** [LOOK (=SCAN but just seeks until no more requests in that direction)]
  - 11000, 20000, 30000, 6000, 9000 [Circular-LOOK (or C-SCAN): seek to lowest track (or start)]
  - 6000, 11000, 30000, 9000, 20000 [First come, first served]

**Questions 20-22**

- The NOOP disk scheduler in the Linux 2.6 kernel is best suited for:
  - flash memory.**
  - database systems.
  - real-time systems.
  - I/O-intensive systems.
- In the POSIX VFS architecture, which object do you need to access to remove a file?
  - superblock
  - inode**
  - dentry
  - File

In-class hint: or change permissions; or change the owner
- The File Allocation Table of Microsoft's FAT32 file system is an example of:
  - linked allocation.**
  - contiguous allocation.
  - indexed allocation.
  - combined indexing.

In-class hint: derivative, not a direct example

**Questions 23-25**

- Larger cluster sizes DO NOT:
  - improve file data access performance. [Usually yes: we read more contiguous data]
  - reduce external fragmentation.** [Not an issue when allocation units are the same size]
  - reduce file fragmentation. [A file has more contiguous chunks so it has less opportunity to get fragmented]
  - increase internal fragmentation. [larger allocation units INCREASE internal fragmentation]
- Linux's ext2, ext3, and ext4 file systems DO NOT use block groups to:
  - keep a file's blocks close to each other.
  - keep an i-node in close proximity to the file.
  - provide a certain degree of fault resiliency: the file system is not destroyed if one group is destroyed.
  - make it easy to locate free blocks.**
- Which form of journaling has the least risk of data corruption?
  - ordered journaling [write data; then just journals metadata]
  - full data journaling** [journal data and metadata]
  - writeback journaling [metadata journaling but no ordering of data blocks first]
  - asynchronous journaling [huh?]

**Questions 26-27**

- Recovery after a crash in a journaled file system involves:
  - replaying all properly terminated transactions in the log (those marked with a transaction end).**
  - replaying all transactions in the log, including partially-logged ones.
  - marking all transactions in the log as invalid and clearing the log.
  - checking the file system integrity.
- The JFFS2 file system performs
  - static wear leveling [we move static data to used blocks to distribute the wear]
  - dynamic wear leveling**
  - both static and dynamic wear leveling
  - no wear leveling but it is optimized for the longer write times of flash media