

Section 6: I/O and File Systems

Jason Jia

March 17, 2015

Contents

1	Warmup	2
1.1	True/False	2
2	Vocabulary	3
3	Problems	4
3.1	Clock Algorithm	4
3.2	Disabling Interrupts	4
3.3	File Systems	4
3.4	Disks	6

1 Warmup

1.1 True/False

Write True or False for each question and justify your answer.

1. If a particular IO device implements a blocking interface, then you will need multiple threads to have concurrent operations which use that device.

2. For IO devices which receive new data very frequently, it is more efficient to interrupt the CPU than to have the CPU poll the device.

3. With SSDs, writing data is straightforward and fast, whereas reading data is complex and slow.

4. User applications have to deal with the notion of file blocks, whereas operating systems deal with the finer grained notion of disk sectors.

5. Directories in UNIX are basically files with pointers to inodes in them.

2 Vocabulary

- **Second Chance Algorithm** A modified form of FIFO that is used to approximate LRU. Each page also has a reference of a use bit to keep track of whether that page has been accessed. It works by looking at the front of the queue as FIFO does, but instead of immediately paging out that page, it checks to see if its referenced bit is set. If it is not set, the page is swapped out. Otherwise, the referenced bit is cleared, the page is inserted at the back of the queue and the process is repeated until a page is swapped out.
- **Clock Algorithm** Clock is a more efficient version of FIFO than Second-chance because pages don't have to be constantly pushed to the back of the list. The clock algorithm keeps a circular list of pages in memory, with the "hand" pointing to the last examined page in the list. When a page fault occurs and no empty frames exist, then the use bit is inspected at the hand's location. If U is 0, the new page is put in place of the page the "hand" points to, otherwise the U bit is cleared. Then, the clock hand is incremented and the process is repeated until a page is replaced.
- **I/O** In the context of operating systems, input/output (I/O) consists of the processes by which the operating system receives and transmits data to connected devices.
- **Controller** The operating system performs the actual I/O operations by communicating with a device controller, which contains addressable memory and registers for communicating with the CPU, and an interface for communicating with the underlying hardware. Communication may be done via programmed I/O, transferring data through registers, or Direct Memory Access, which allows the controller to write directly to memory.
- **Interrupt** One method of notifying the operating system of a pending I/O operation is to send an interrupt, causing an interrupt handler for that event to be run. This requires a lot of overhead, but is suitable for handling sporadic, infrequent events.
- **Polling** Another method of notifying the operating system of a pending I/O operation is simply to have the operating system check regularly if there are any input events. This requires less overhead, and is suitable for regular events, such as mouse input.
- **Response Time** Response time measures the time between a requested I/O operation and its completion, and is an important metric for determining the performance of an I/O device.
- **Throughput** Another important metric is throughput, which measures the rate at which operations are performed over time.
- **Asynchronous I/O** For I/O operations, we can have the requesting process sleep until the operation is complete, or have the call return immediately and have the process continue execution and later notify the process when the operation is complete.
- **inode** In UNIX based operating systems, an inode is a data structure used to represent a filesystem object. It contains attributes about the file, as well as the locations of the disk blocks where the file contents reside. For large files, it also contains pointers to multi-level disk block locations.

3 Problems

3.1 Clock Algorithm

Suppose that we have a 32-bit virtual address split as follows:

10 Bits	10 Bits	12 Bits
Table ID	Page ID	Offset

Show the format of a PTE complete with bits required to support the clock algorithm.

For this problem, assume that physical memory can hold at most four pages. What pages remain in memory at the end of the following sequence of page table operations and what are the use bits set to for each of these pages:

- Page A is paged in
- Page B is paged in
- Page C is paged in
- Page A is accessed
- Page C is accessed
- Page D is paged in
- Page B is accessed
- Page D is accessed
- Page A is accessed
- Page E is paged in
- Page F is paged in

3.2 Disabling Interrupts

We looked at disabling CPU interrupts as a simple way to create a critical section in the kernel. Name a drawback of this approach when it comes to I/O devices.

3.3 File Systems

Indirect blocks look a lot like another data structure we looked at earlier this semester. What is the other data structure? Can you compare the tradeoffs you make in either case?

How many times is the disk accessed when you type `ls /home/cs162`? What is each access for?

Assume you already know the block number of the root inode and that `ls` and all of its corresponding libraries are already loaded in memory.

Assuming the total response time to read an inode or block from disk is 5ms, and all inodes and directories consume one block, how long does this take?

Now say the following is true:

- The root directory listing is cached in memory.
- The disk controller is using a track buffer.
- I-nodes are always stored on the same cylinder as the blocks they refer to.

Assume that data in the track buffer or in memory is free (e.g. 0ms) to access. Now much time would it take to read the contents of `/home/cs162`?

3.4 Disks

What are the major components of disk latency? Explain each one.

In class we said that the operating system deals with bad or corrupted sectors. Some disk controllers magically hide failing sectors and re-map to back-up locations on disk when a sector fails.

If you had to choose where to lay out these back-up sectors on disk - where would you put them? Why?

How do you think that the disk controller can check whether a sector has gone bad?

Can you think of any drawbacks of hiding errors like this from the operating system?