

B. How path look up occurs across i-nodes and datablocks?

- When a file opened, the file system must take the file name supplied and locate its disk block i.e.
- Path name /usr/ast/mbox is looked up.
 - File system locates the root directory
 - Reads the root directory and looks first component of the path, i.e. /usr, in the root directory to find i-node number of the file /usr
 - Since each i-node has a fixed location on the disk, it is straightforward
 - From this i-node, the system locates the directory for /usr and looks up for the next component, ast, in it
 - Keep repeating until reaching mbox
 - i-node for this file file is read into memory and kept there until the file closed.

- No problem with this question as far as I am concerned
- Question asked with respect to 5.1.2 algorithm
- Variables with respect to the formula
- Suppose now we have two objects: object one already allocated and object two

Root directory	
1	.
1	..
4	bin
7	dev
14	lib
9	etc
6	usr
8	tmp

Looking up
USR will be
i-node 6

i-node 6
is for usr/
directories

• Mode
• size
• times

132

i-node 6
says that
/usr is in
block 132

Block 132
is /usr/
directories

6.
..
19.

dict
eric

31.
26.

45.
bol

/usr/ast

is i-node 26

i-node 26
is for
/usr/ast/
directories

• Mode
• size
• times

406

grants
books

60. mbox

91. minix

17. src

/usr/ast/mbox

is i-node 60

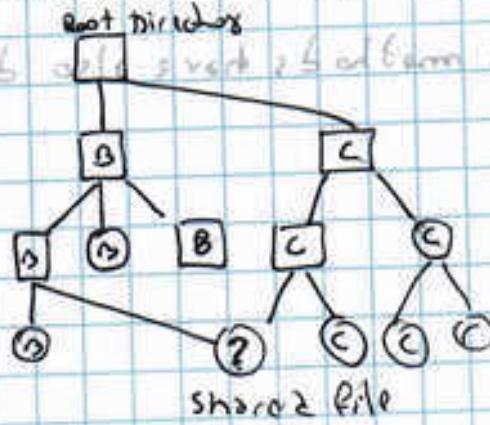
(Steps into in looking up /usr/ast/mbox)

Note: Relative path names are looked up the same way as absolute ones, only starting from the working directory instead of starting from the root directory

9. How shared files work?

- When a file is shared so it can appear simultaneously in different directories, the file system became a Directed Acyclic Graph (DAG), rather than a tree.

For example: lets assume one file belonging to user C is shared with user B by making a connection, link:



Final review

iii. Disadvantage:

I. If each i-node has room for a fixed number of disk addresses, what happens when a file grows beyond limit?

- Solution: Reserve the last disk address not for data block, but instead for the address of a block containing more disk block addresses.

f. UNIX i-node structure:

i. i-nodes contain attributes such as:

I. file size

II. creation, last access, and last modified times

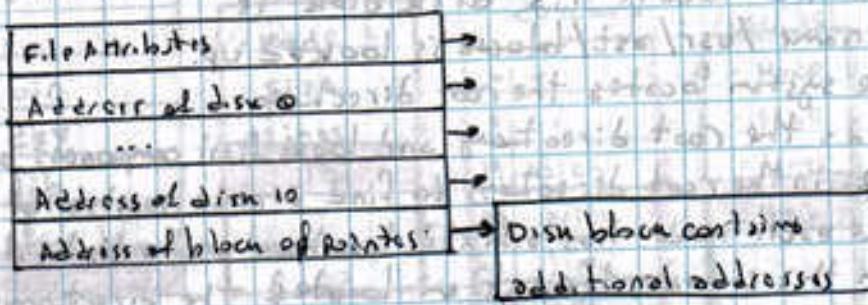
III. protection information

IV. count of the # of directories entries that point to the i-node, needed due the links

A. When New link is made to an i-node, count increased

B. When a link is removed, count is decremented. When it gets to 0, the i-node is reclaimed and the disk blocks put back in the freelist

ii. Keeping track of disk blocks is done using generalization.



I. The first 10 disk addresses are stored in the i-node. For small files, which are in memory, is fetched from disk to main memory when file is opened.

II. For large file, one of the addresses in the i-node is the address of a disk block called a single indirect block. This block contains additional additional disk addresses.

A. If not enough double and triple indirect block can be used as needed.

g. How shared files work?

158 b. Advantages: Shared file to appear simultaneously in different directories (even if these directories belong to different users)

c. disadvantages:

I. if directory really do contain disk addresses, then

- copy of the disk addresses will have to be made in B's directory when the file is linked

II. if either B or C subsequently appears to the file, the new blocks will be listed only in the directory of the user doing the append. The changes will not be visible to the other user, thus defeating the purpose of sharing.

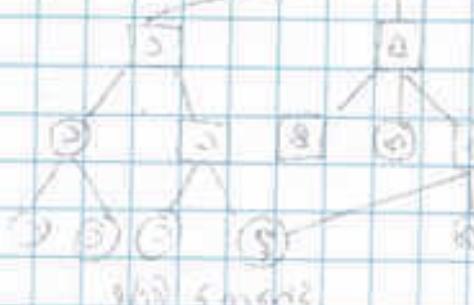
d. Solutions to disadvantages:

I. Disk blocks are not listed in directories, but in a little data structure associated with the file itself.

The directories would then point to the little data structure (i-node).

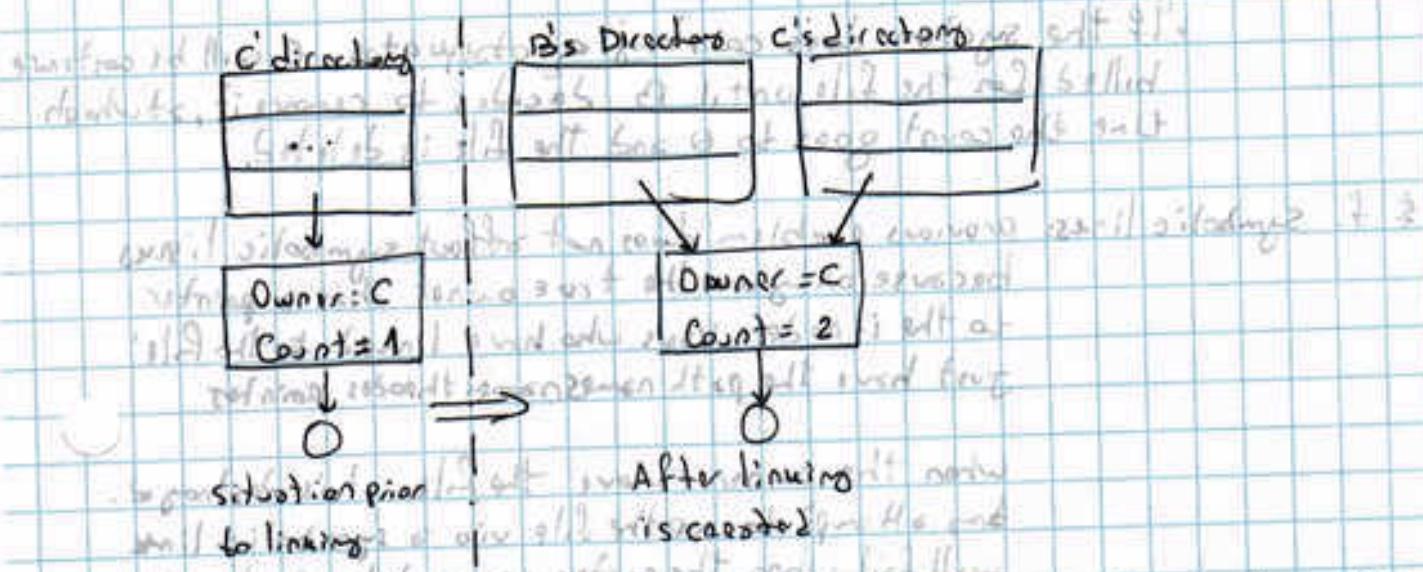
II. B links to one of C's files by having the system create a newfile, of type LINK, and entering that file in B's directory. The newfile contains just the path name of the file to which it is linked. When B reads from the linked file, the OS sees that the file being read from is of type LINK, looks up the name of the file, and reads that file (this approach is called symbolic linking, to contrast it with traditional hard linking).

III. These methods have also drawbacks

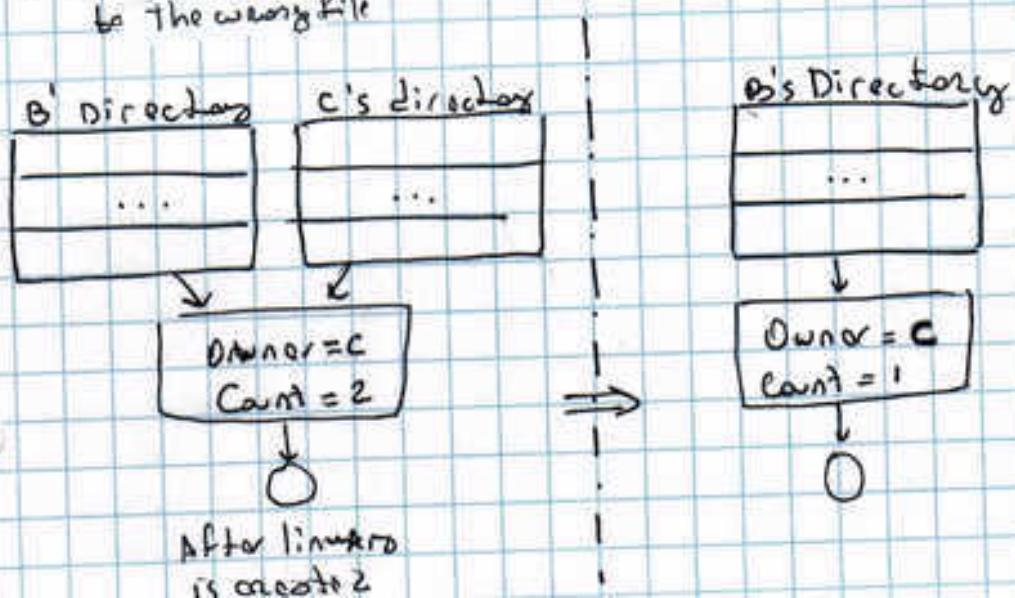


- e. disadvantages of the solutions

I. For the first solution: the moment that B links to the shared file, the i-node records the file's owner as C. Creating a link doesn't change the ownership, but does increase the link count in the inode, so the system knows how many directory entries currently point to the file.



- A. If C subsequently tries to remove the file, the system is faced with a problem.
- If it removes the file and clear the i-node, B will have a directory entry pointing to an invalid i-node.
 - If the i-node is later reassigned to another file, B's link will point to the wrong file.



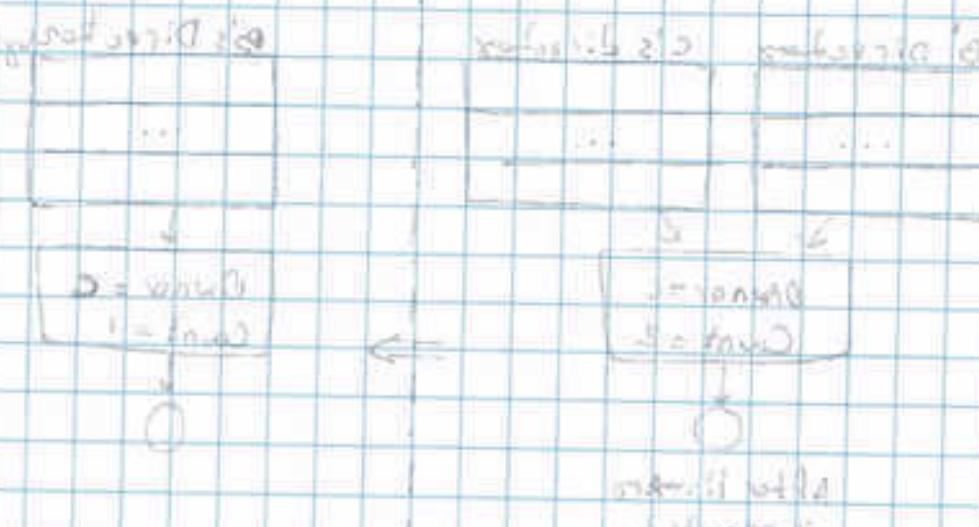
- 160 c. System can see from the count in the i-node that the file is still in use, but there is no easy way for it to find all the directory entries for the file, in order to erase them.
- d. Pointers to the directories cannot be stored in the inode because there can be an unlimited number of directories.
- e. Only thing to do: Remove C's directory entry, but leave the i-node intact, with count set to 1
- Now B is only user having a directory entry for file owned by C.
- If the system does accounting or quotas, C will be continue billed for the file until B decides to remove it, at which time the count goes to 0 and the file is deleted.

f. Symbolic links: previous problem does not affect symbolic links because only one the true owner has a pointer to the i-node. Users who have linked to the file just have the path names; not i-nodes pointers

when the owner removes the file, it is destroyed.

Any attempt to use the file via a symbolic link will fail when the system is unable to locate the file. Removing a symbolic link does not affect the file.

I. Problem: extra overhead required to follow after



10. Disk Free Space (block) management

- a. Using a linked list of disk blocks, with each block holding as many free disk block numbers as will fit.

- I. with 1KB blocks, 32-bit disk block number, each block on the free list holds the number of 255 free blocks
 ie: $500 \text{ GB HD} = 488 \cdot 10^6 \text{ disk blocks}$
 to store all these address at 255 per block = $1.9 \cdot 10^6$ blocks

Given II. Free disk blocks: 16, 17, 18

42	230	36
136	162	238
210	612	897
97	942	422
48	214	140
63	160	223
21	664	160
262	216	126
310	320	142
516	180	141

1 KB Disk Block ≈ ... ≈
 =
 256 32-bit Disk Block numbers

- b. Using a bitmap, a disk with n blocks requires a bitmap with n bits. Free blocks are represented by 1s in the map, allocated blocks by 0s.
 ie: $500 \text{ GB} = 488,000,000 \text{ bits}$ for map
 $= 60,000 \text{ 1KB blocks to store}$

I. Therefore bitmaps require less space than linked list

II.

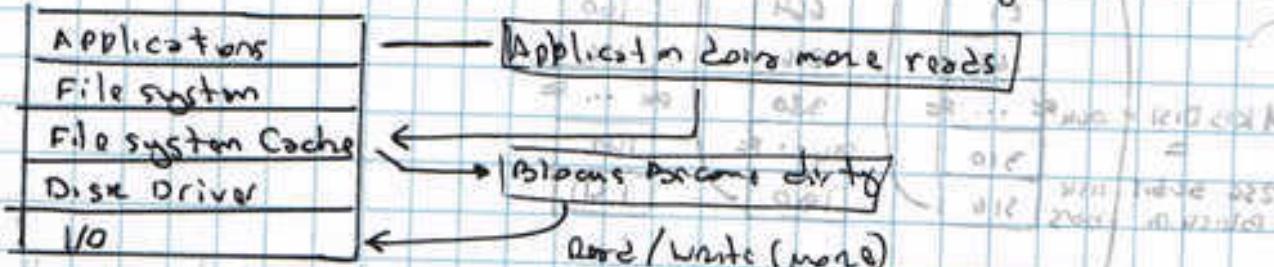
1	0	0	1	1	0	1	1	0	1	1	0
0	1	1	0	1	1	0	1	1	0	1	0
1	0	1	0	1	0	1	1	0	1	1	1
..											
1	0	1	1	1	1	0	0	0	1	1	1

11. File System Cache

written by [unclear]

- a. block cache or buffer cache is a technique used to reduce disk accesses.
- b. A cache is a collection of blocks that logically belong on the disk but are being kept in memory for performance reasons.
- c. Small area in main memory that stores frequently accessed data blocks in the file system.
- d. Before accessing the disk, look in the filesystem cache.
 - I. if you find it in the file system cache, there is no need to go to the disk.
- e. periodically, purge the cache of infrequently used data blocks.

claim: If the cache works well, then most I/O accesses to the physical disk will be writes. why?



12. Data Structure for File-System Cache

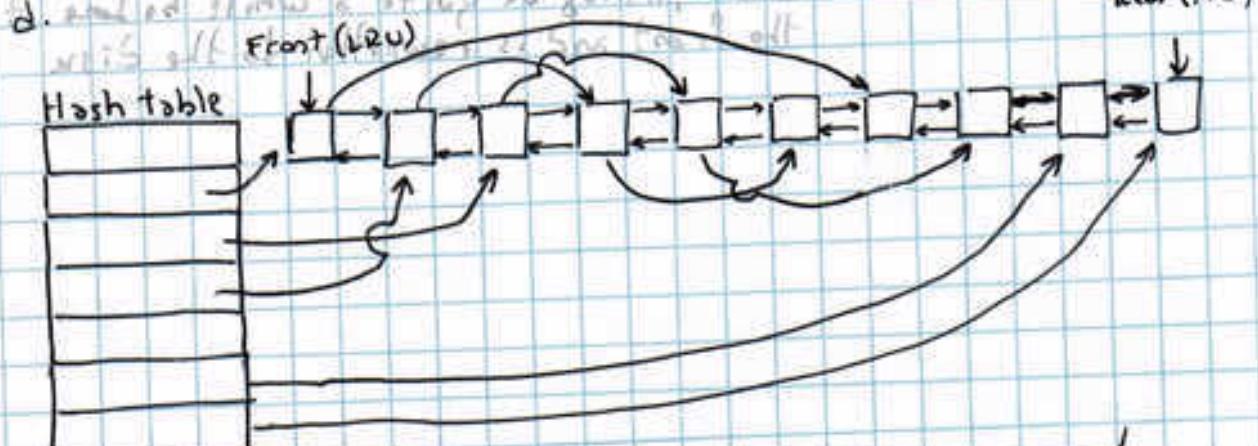
a. since many blocks in the cache, there is need of some method to determine quickly if a given block is present.

b. A method usually used is to hash the device and disk address and look up the result in a hash table.

i. All the blocks with the same hash value are chained together on a linked list so that the collision chain can be followed.

ii. When a block has to be loaded into a full cache, some block has to be removed (and rewritten to the disk if it has been modified since being brought)

c. Similar to Page replacement algorithm. Normally the file system cache and page system go together in Linux systems since they have the same size as page system even though they don't have to.



Another doubly-linked list maintained to identify least recently used (LRU) pages that are periodically purged to disk

Pages in cache are normally looked up via a hash table for fast access. Best case O(1); worst case O(n)

i. Similar to paging and page replacement algorithm such as F, MRU, second chance and LRU are applicable; however the difference is that cache references are relatively infrequent, so it is feasible to keep all the blocks in exact LRU order with linked list after each access since it requires no power in placing information at bottom.

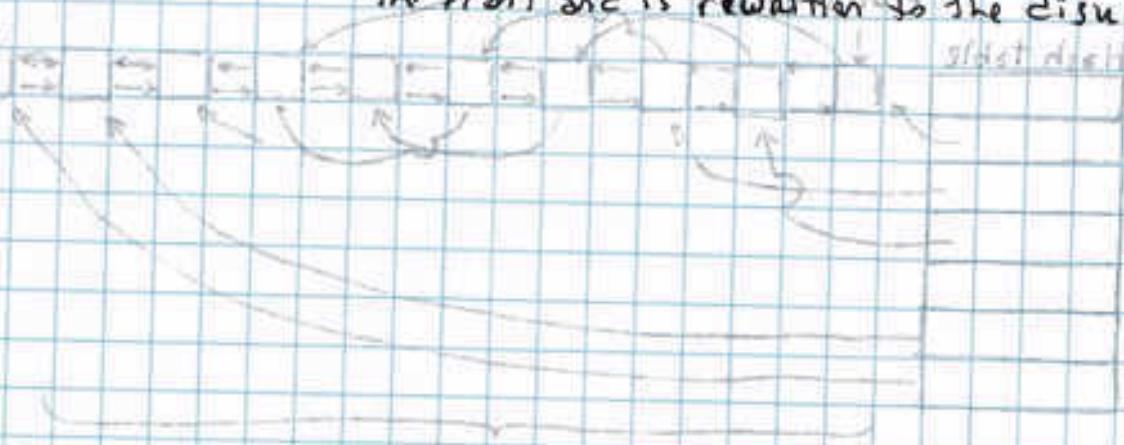
ii. In addition to collision chains starting at the hash table, there is a bidirectional list running through all the blocks in the order of usage, with the least recently used block at the end of the list.

iii. When a block is referenced, it can be moved from its position on the bidirectional list and put at the end. In this way exact LRU order can be maintained.

iii. Catch: situations in which exact LRU possible, it turns out LRU is undesired because it causes inconsistency.

I. If a control block such as an i-node block, which will not be modified, but not rewritten to the disk, a crash will leave the filesystem in an inconsistent state.

II. If the i-node block is put at the end of the LRU chain, it may be quite a while before it reaches the front and is rewritten to the disk.



The following diagram illustrates the linked list structure:

Ques. Given a linked list with nodes x and y ,
 (i) O(1) insertion; (ii) O(n) insertion for n nodes.

13. Performance impact of i-node placement

- a. Apart of caching and read ahead, reducing the amount of disk arm motion by putting blocks access & in sequence closer to each other (preferably in the same cylinder).
- i. When output file is written, FS tries to allocate these blocks one at a time, on demand.

ii. If free blocks are recorded in bitmaps in main memory, it's easy to choose free block as close to previous block.

b. In free list, some block clustering can be done by keeping track of the disk storage not in blocks, but in a group of consecutive blocks.

- c. Another performance bottleneck in systems using i-nodes is that reading even a short file requires two disk access: one for the i-node, and one for the block. As of the usual i-node placement:

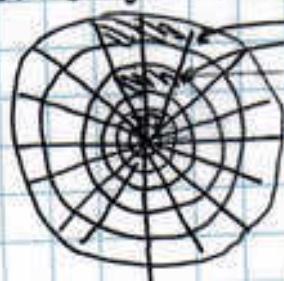


i. here all i-nodes are near the beginning of

the disk, so the average distance between an i-node and its block will be half the number of cylinders, requiring long seeks.

- c. Another performance is to put the i-nodes in the middle of the disk, rather than at the start, thus reducing average seek between the i-node and the first block by a factor of two.

- d. Another idea is to divide the disk into cylinder groups, each with its own i-nodes, blocks, and free list (McKusick 1984). When creating a new file, any i-node can be chosen, but an attempt is made to find a block in the same cylinder group as the i-node. If none is available, then block in a nearby cylinder group is used.



Disk is divided into cylinder groups, each with its own i-nodes

cylinder group

14. Log-Structured File Systems (LFS)

- a. With CPUs faster & memory larger
 - i. disk caches can also be larger
 - ii. increasing number of read requests can be come from cache with no disk access needed
 - iii. Most disk accesses will be writes
 - iv. read ahead no longer needed to gain performance
 - v. However, small writes are inefficient
 - vi. delayed writes expose fs to inconsistency problems therefore i-nodes writers are done immediately

b. LFS strategy: structure entire disk as a log

- i. All writes initially buffered in memory are collected into a single segment and written to the disk as a single continuous segment at the end of the log.

ii. When file opened, locate i-nodes, then find blocks

iii. If file is overwritten, its i-node will now point to the new blocks, but old ones will still be occupying space in previous log segment

I. Solution: LFS has a cleanup thread that scans the log to compact it

Input/Output Subsystem1. I/O devices categories

a. block devices: store information in fixed-size blocks, each one with its own address.

i. All transfers are in units of one or more entire (consecutive) blocks

ii. It is possible to read or write each block independently of all the other ones

iii. i.e.: Hard Disk, CD ROMs

Note: sometimes boundary between devices that are blockwise addressable and those that are not is not well defined.

b. Character device: A character device delivers or accepts a stream of characters, without regards to any block structure.

i. It is not addressable and does not have seek operation

ii. i.e.: Printer, mice,

c. Some devices are not block or characters such as Clocks.

i. Clocks are not accessible but can cause interrupt at well-defined intervals

ii. do not generate or accept character streams

iii. They do cause interrupts at well-defined intervals

2. Device Controllers

a. I/O devices have components: mechanical and electronic components

b. The electronic component is the device controller which may be able to handle multiple devices. Electronic components are called device controllers or adapter.



168 C. Controller tasks

Driver level

- i. Intermediary between I/O devices and CPU
(commands from the CPU to the I/O device back and forwards)

- ii. Convert serial bit stream to block of bytes and perform error correction as necessary

- iii. Block of bytes is first assembled, bit by bit, in a buffer inside the controller. After its checksum has been verified and the block has been declared to be free of errors, it can then be copied to main memory and made available.

3. Memory-Mapped I/O

- a. Each controller has few registers that are used for communicating with the CPU.

- i. By writing into these registers, the OS command the device to deliver data, switch itself on/off, or other action.

- ii. By reading these registers, the OS can know what is the state of the device, whether is prepared to accept a new command, and so on.

- b. In addition to control registers, many devices have a data buffer that the OS can write & read.

- c. How the CPU communicates with the control registers and the device data buffer?

- i. Each control register is assigned an I/O port number

I. set of all the I/O ports form the I/O port space

II. the I/O port space is protected so user programs cannot access, only the OS

III. Using a special I/O instruction the CPU read in control register port and store the result in the CPU register

0xFFFF - Also the CPU can write the content of CPU registers to the control register



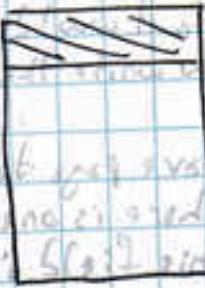
Final review

ii. Another way is to map all the control registers into the memory space.

I. Each control register is assigned a unique memory address to which no memory is assigned. This system is called memory-mapped I/O.

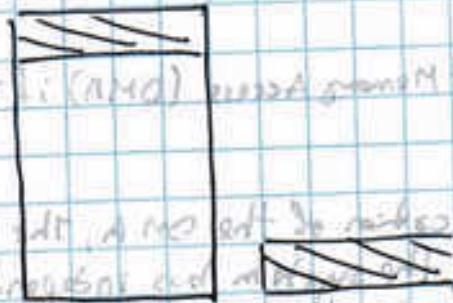
II. Assigned addresses are at the top of the address space.

one address space



iii. Hybrid scheme, with memory-mapped I/O data buffers & separate I/O ports for the control registers

two addresses



I. When the CPU wants to read a word, either from memory or from an I/O port, it puts the address it needs on the bus' address lines and then asserts a READ signal on a bus' control line.

II. A second signal line is used to tell whether I/O space or memory space needs.

A. If it is memory space, the memory responds to the request.

B. If it is I/O space, the I/O device responds to the request.



This mem port is to allow I/O devices access to memory

170 c. If there is only memory space, every memory module and every I/O device compares the address lines to the range of addresses that it services.

• If the address falls in its range, it responds to the request.

III. Since no address is ever assigned to both memory and an I/O device, there is no ambiguity and no conflict.

iv. How if the same way you read and write to memory, you could do the same with devices?

two Address



I. The first part of memory it would be reserved for the I/O device, so if you write there, memory mapped

II. In page tables, we have page table entries.

A. In the page entries there is one field that tells care of the I/O. This field is the disable cache (or don't cache).

B. Every read and write will be performed directly to the device when disable cache tells.

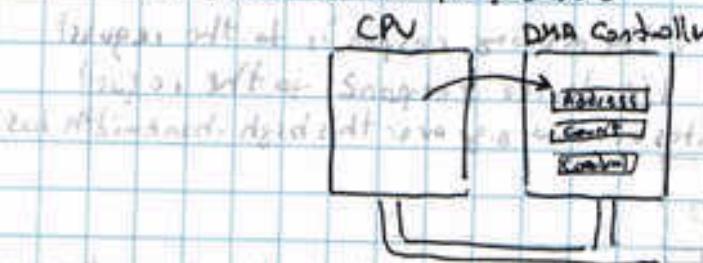
4- Direct Memory Access (DMA)

a. The OS can only use Direct Memory Access (DMA) if the Hardware has a DMA controller.

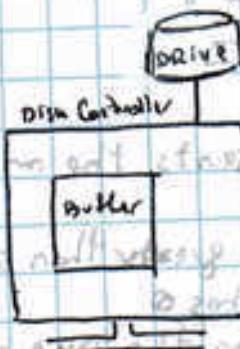
b. No matter the physical location of the DMA, the DMA controller must be able to access the system bus independent of the CPU.

I. These include memory address registers, a byte count register, and one or more control registers. (to specify I/O port base, direction of transfer (memory to I/O device), transfer unit (byte or word), and number of bytes to transfer in one burst)

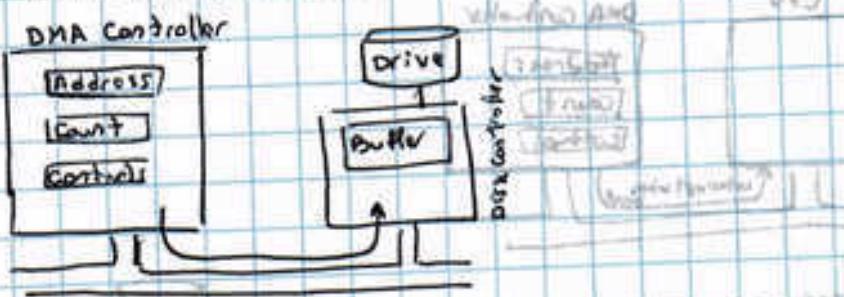
c. The CPU programs the DMA controller by setting its registers so it knows what to transfer.



- Final review: what does DMA do, what's its function, what's its advantage?
- c. The CPU also issues a command to the disk controller telling it to read data from the disk into its internal buffer and verifying the checksum. When valid data are in the disk controller's buffer, DMA can begin.

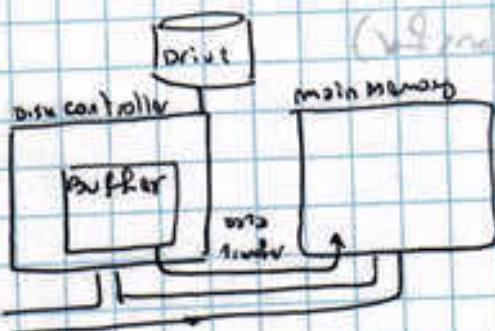


- e. The DMA controller initiates the transfer by issuing a read request over the bus to the disk controller.

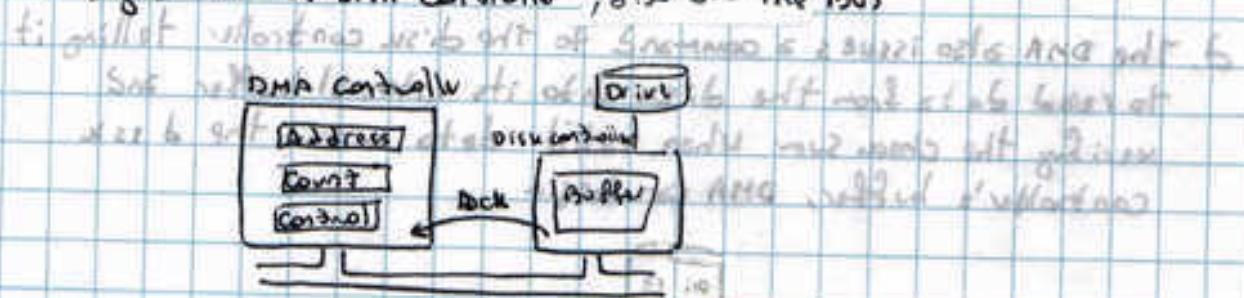


The disk controller read request as any other request since doesn't know or care if request came from the CPU or DMA controller.

- f. The memory address to write to is on the bus' address lines so when the disk controller fetches the next word from its internal buffer, it knows where to write it. The write to memory is another standard bus cycle.

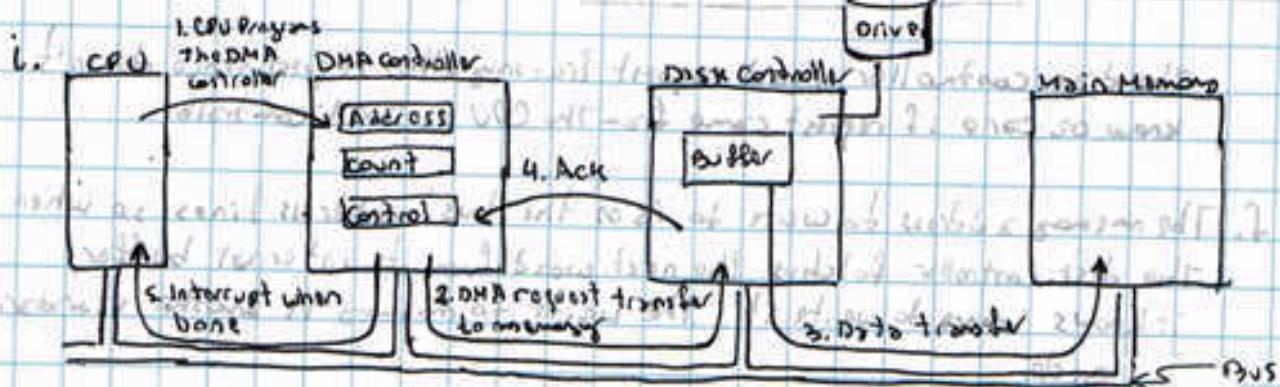
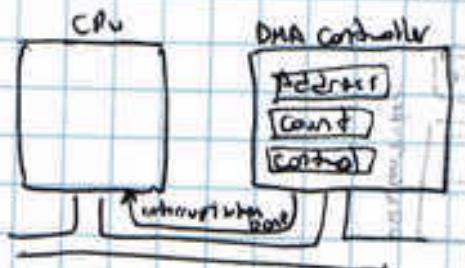


172 g. When the write is complete, the disk controller sends an acknowledgement signal to the DMA controller, also over the bus



h. the DMA controller then increments the memory address to use and decrement the byte count

- if the byte count is still greater than 0, repeat previous process of transfer until count reaches 0
- If byte count reaches 0, then the DMA controller interrupts the CPU to let it know that the transfer is now complete

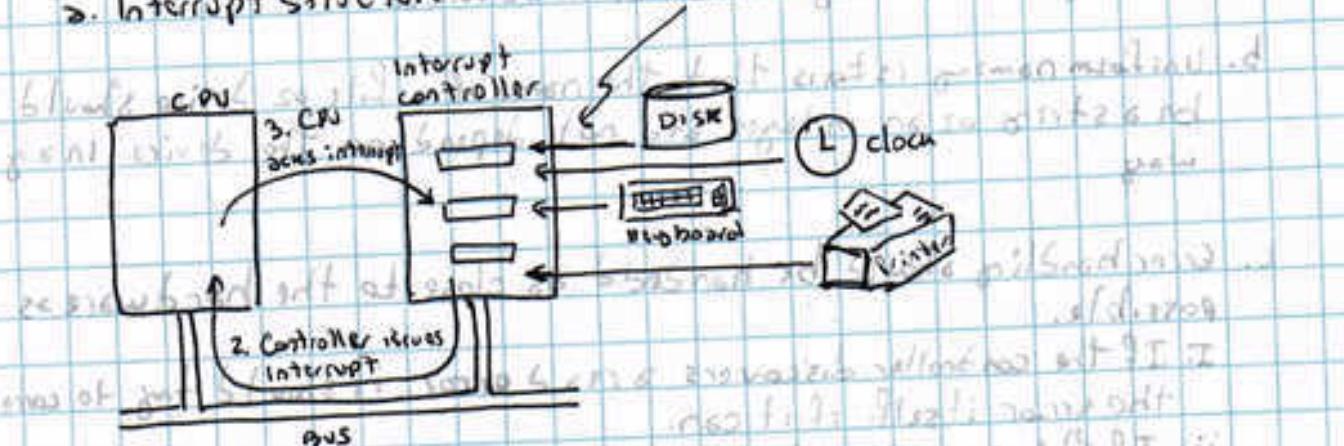


(Operation of a DMA transfer)

5. Interrupts Revised

a. Interrupt Structure

- 1. Device finished its job



- 2. Controller issues interrupt
 - 3. CPU gets interrupt
- Interrupt controller will decide which devices can generate IRQ and what IRQ they can generate
 - If interruption controller do not recognize the IRQ, the IRQ will be lost
 - Interrupt controller designate IRQs to the devices
 - When an I/O device has finished the work given to it, it cause an interrupt (assuming that interrupts have been enabled by the OS). It does this by asserting a signal on a bus line that it has been assigned
 - This signal is detected by the interrupt controller chip on the motherboard, which then decides what to do.
 - If no other interrupt are pending, the interrupt controller processes the interrupt immediately.
 - If another one is in process, or another device has made a simultaneous request on a higher-priority interrupt request line on the bus, the device is just ignored for the moment. In this case it continues to assert an interrupt signal on the bus until it is serviced by the CPU.