

CS 3204 Operating Systems

Lecture 23
Godmar Back

Announcements

- Project 2 Grades on Curator
- Reading assignment:
 - Chapter 10 + 11 + 12
- Project 4 Help Sessions
 - (1) Tonight: Thursday Nov 9: 7pm McB 126
 - (2) Monday night: Nov 13: 7pm Norris 306

Filesystems

Files vs Disks

File Abstraction

- Byte oriented
- Names
- Access protection
- Consistency guarantees

Disk Abstraction

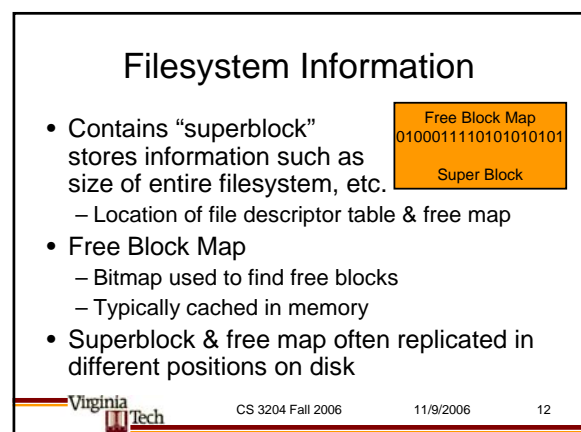
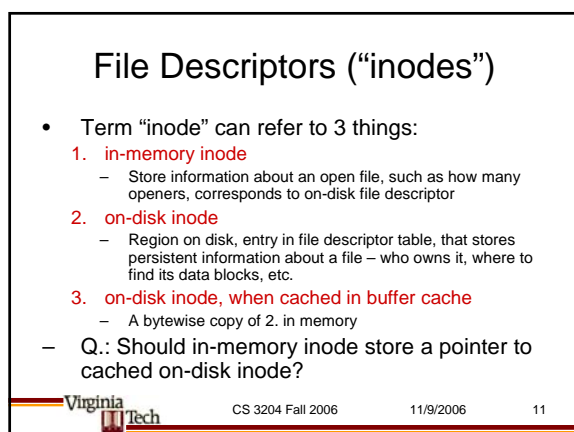
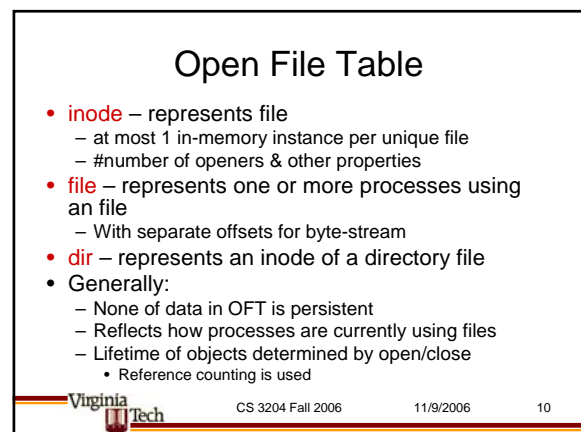
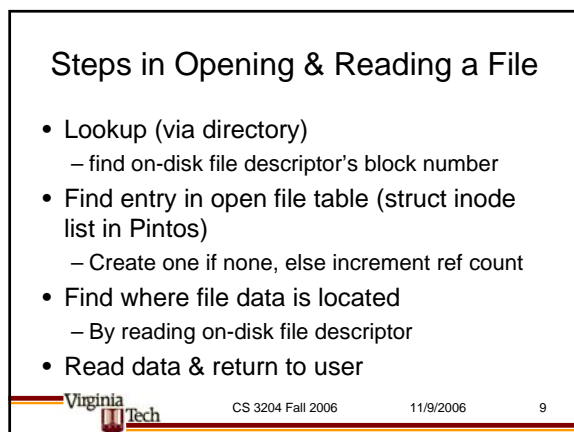
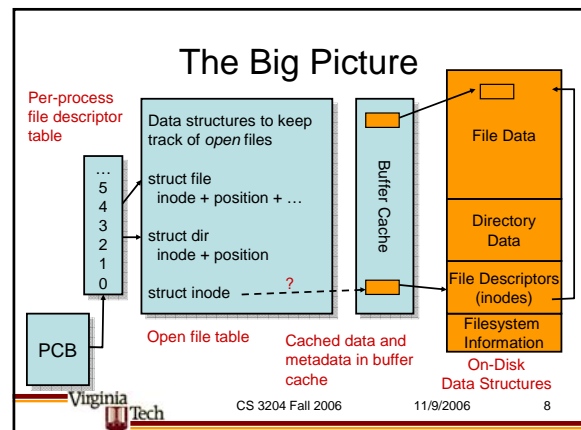
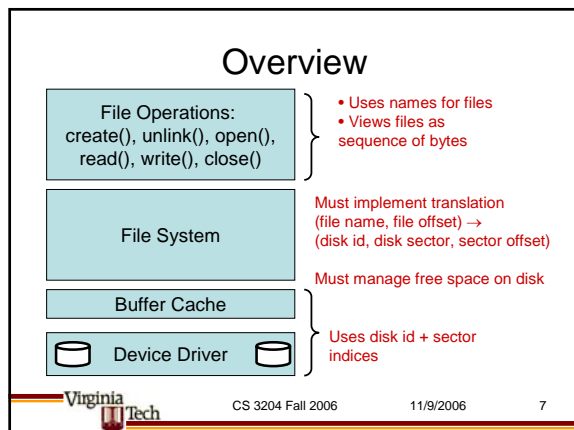
- Block oriented
- Block #s
- No protection
- No guarantees beyond block write

Filesystem Requirements

- Naming
 - Should be flexible, e.g., allow multiple names for same files
 - Support hierarchy for easy of use
- Persistence
 - Want to be sure data has been written to disk in case crash occurs
- Sharing/Protection
 - Want to restrict who has access to files
 - Want to share files with other users

FS Requirements (cont'd)

- Speed & Efficiency for different access patterns
 - Sequential access
 - Random access
 - Sequential is most common & Random next
 - Other pattern is Keyed access (not usually provided by OS)
- Minimum Space Overhead
 - Disk space needed to store metadata is lost for user data
- Twist: all metadata that is required to do translation must be stored on disk
 - Translation scheme should minimize number of additional accesses for a given access pattern
 - Harder than, say page tables where we assumed page tables themselves are not subject to paging!



File Allocation Strategies

- Contiguous allocation
- Linked files
- Indexed files
- Multi-level indexed files

Contiguous Allocation



- Idea: allocate files in contiguous blocks
- File Descriptor = (first block, length)
- Good sequential & random access
- Problems:
 - hard to extend files – may require expensive compaction
 - external fragmentation
 - analogous to segmentation-based VM
- Pintos's baseline implementation does this

Linked Files



- Idea: implement linked list
 - either with variable sized blocks
 - or fixed sized blocks ("clusters")
- Solves fragmentation problem, but now
 - need lots of seeks for sequential accesses and random accesses
 - unreliable: lost first block, lose file
- Solution: keep linked list in memory
 - DOS: FAT File Allocation Table


DOS FAT

- FAT stored at beginning of disk & replicated for redundancy
- FAT cached in memory
- Size: n-bit entries, m-bit blocks → $2^{(m+n)}$ limit
 - n=12, 16, 28
 - m=9 ... 15 (0.5KB-32KB)
- As disk size grows, m & n must grow
 - Growth of n means larger in-memory table

Filename	Length	First Block
"a"	2	1
"b"	4	3
"c"	3	12
"d"	1	4


1	6
2	0
3	5
4	-1
5	7
6	-1
7	11
8	0
9	-1
10	9
11	-1
12	10

Blocksize Trade-Offs

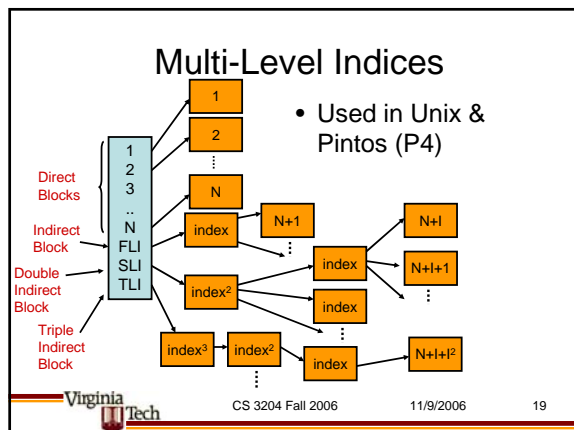


- Assume all files are 2KB in size (observed median filesize is about 2KB)
 - Larger blocks: faster reads (because seeks are amortized & more bytes per transfer)
 - More wastage (2KB file in 32KB block means 15/16th are unused)
- Source: Tanenbaum, *Modern Operating Systems*

Indexed Allocation



- Single-index: specify maximum filesize, create index array, then note blocks in index
 - Random access ok – one translation step
 - Sequential access requires more seeks – depending on contiguous allocation
- Drawback: hard to grow beyond maximum



Multi-Level Indices

- If $\text{filesize} < N * \text{BLKSIZE}$, can store all information in direct block array
 - Biased in favor of small files (ok because most files are small...)
- Assume index block stores I entries
 - If $\text{filesize} < (I + N) * \text{BLKSIZE}$, 1 indirect block suffices
- Q.: What's the maximum size before we need triple-indirect block?
- Q.: What's the per-file overhead (best case, worst case?)

Virginia Tech CS 3204 Fall 2006 11/9/2006 20

