

# More on File System

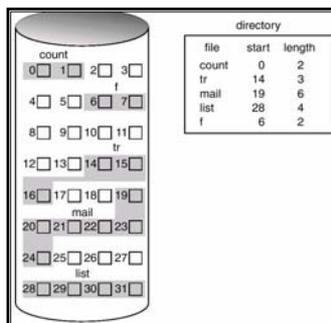
CS 256/456

Dept. of Computer Science, University of Rochester

## File on the Disk

- Disk basic allocation unit is a block (e.g., 512 bytes)
- File system may choose to use a larger block size (e.g., 4KB)
- File allocation methods
  - How disk blocks are allocated for files
  - Contiguous allocation, linked allocation, indexed allocation
  - Metrics: access speed (sequential & random), space utilization

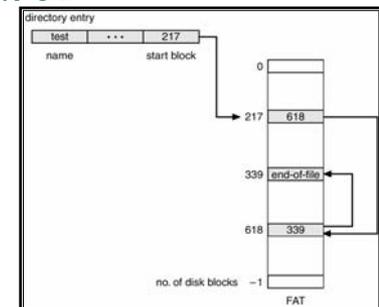
## Contiguous File Allocation



- Each file occupies a set of contiguous blocks on the disk.
- Advantage:
  - Simple - only starting location (block #) and length (number of blocks) are required.
  - Fast sequential; also quite fast random access.
- Disadvantage:
  - External fragmentation
  - Inflexible when appending to a file

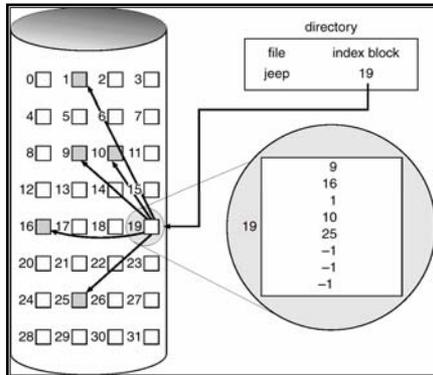
## Linked File Allocation

- Each file is a linked list of disk blocks
  - each block contains a next pointer
  - directory only needs to store the pointer to the first block
  - blocks may be scattered anywhere on the disk.
- Advantage
  - Space efficient
  - Flexible in appending
- Disadvantage:
  - Poor access speed (sequential? random?)



## Indexed File Allocation

- Brings all pointers together into the *index block*.

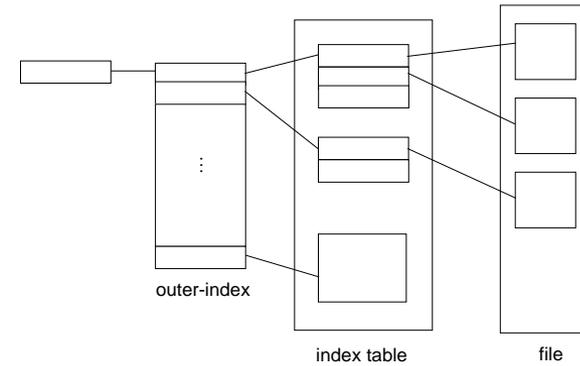


3/21/2007

CSC 256/456 - Spring 2007

5

## Multi-level Indexed File Allocation

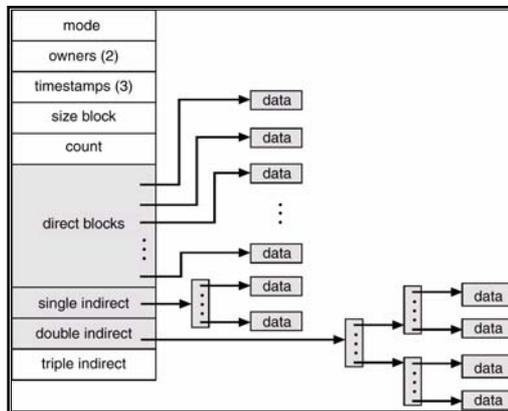


3/21/2007

CSC 256/456 - Spring 2007

6

## UNIX (4K bytes per block)



3/21/2007

CSC 256/456 - Spring 2007

7

## Directory on the Disk

- Directory is a container of files
- For space management, similar to files
- But for directory, file system does care its content
  - Linear list of file names and attributes (including pointers to the data blocks)
    - time-consuming to search an item
  - Hash Table - using a link list to chain all files hashed to the same value
    - Pro: decreases directory search time
    - Con: increased complexity, a little waste of space
- how much benefit does it really provide?

3/21/2007

CSC 256/456 - Spring 2007

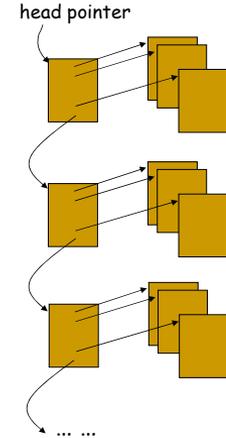
8

## Where to put file attributes?

- File control block - data structure including all attributes for a file
- Where to put the file control block?
  - In directory data structure
    - Hard to share files through links?
  - In system-level dedicated data structure
    - inode

## Free-Space Management

- Free-space management for memory?
- Bit map and linked free block list
- Space overhead?
- Efficiency?
  - getting the address of one free block
  - getting the addresses of a number of free blocks
- Grouping



## File System Caching

- File content is cached in memory buffer for later reuse
  - what is the basic unit of such caching?
- Replacement policy for file system buffer cache
  - LRU replacement is one possibility; but sequential access is very likely in file system I/O
  - MRU or free-behind

## File System Prefetching

- File content is read ahead of time for anticipated use in the near future
- Often sequential (based on past access history on the file)
- What is the advantage of file prefetching?
- What is the danger of file prefetching?
- A balanced scheme that provides a competitive performance to the optimal scheme [Li et al. EuroSys 2007]

## Disclaimer

- Parts of the lecture slides contain original work of Abraham Silberschatz, Peter B. Galvin, Greg Gagne, Andrew S. Tanenbaum, and Gary Nutt. The slides are intended for the sole purpose of instruction of operating systems at the University of Rochester. All copyrighted materials belong to their original owner(s).