

Last Class: Memory Management

- Allocating memory to processes
- Limited physical memory, internal and external fragmentation
- Paging and segmentation
 - Address translation, efficiency
 - Hardware support (e.g., TLB)
- Page replacement algorithms
 - FIFO, MIN, LRU
 - Approximations to LRU: second chance
 - Multiprogramming considerations, thrashing




Today: File System Functionality

Remember the high-level view of the OS as a translator from the user abstraction to the hardware reality.

User Abstraction		Hardware Resource
Processes/Threads	<= OS =>	CPU
Address Space	<= OS =>	Memory
Files	<= OS =>	Disk



File System Abstraction



User Requirements on Data

- **Persistence:** data stays around between jobs, power cycles, crashes
- **Speed:** can get to data quickly
- **Size:** can store lots of data
- **Sharing/Protection:** users can share data where appropriate or keep it private when appropriate
- **Ease of Use:** user can easily find, examine, modify, etc. data



Hardware/OS Features

- Hardware provides:
 - **Persistence:** Disks provide non-volatile memory
 - **Speed:** Speed gained through random access
 - **Size:** Disks keep getting bigger (typical disk on a PC = 500GB)
- OS provides:
 - **Persistence:** redundancy allows recovery from some additional failures
 - **Sharing/Protection:** Unix provides read, write, execute privileges for files
 - **Ease of Use**
 - Associating names with chunks of data (files)
 - Organize large collections of files into directories
 - Transparent mapping of the user's concept of files and directories onto locations on disks
 - Search facility in file systems (Spotlight in Mac OS X)



Files

- **File:** Logical unit of data on a storage device
 - Formally, named collection of related information recorded on secondary storage
 - **Example:** reader.cc, a.out
- Files can contain programs (source, binary) or data
- Files can be structured or unstructured
 - Unix implements files as a series of bytes (unstructured)
 - IBM mainframes implements files as a series of records or objects (structured)
- File attributes: name, type, location, size, protection, creation time



File Operations: Creating a File

- **Create(name)**
 - Allocate disk space (check disk quotas, permissions, etc.)
 - Create a file descriptor for the file including name, location on disk, and all file attributes.
 - Add the file descriptor to the directory that contains the file.
 - Optional file attribute: file type (Word file, executable, etc.)



File Operations: Deleting a File

- **Delete(name)**
 - Find the directory containing the file.
 - Free the disk blocks used by the file.
 - Remove the file descriptor from the directory.

 - Behavior dependent on hard links



File Operations: Open and Close

- **fileId = Open(name, mode)**
 - Check if the file is already open by another process. If not,
 - Find the file.
 - Copy the file descriptor into the system-wide open file table.
 - Check the protection of the file against the requested mode. If not ok, abort
 - Increment the open count.
 - Create an entry in the process's file table pointing to the entry in the system-wide file table. Initialize the current file pointer to the start of the file.
- **Close(fileId)**
 - Remove the entry for the file in the process's file table.
 - Decrement the open count in the system-wide file table.
 - If the open count == 0, remove the entry in the system-wide file table.



OS File Operations: Reading a File

- **Read(fileID, from, size, bufAddress)** - random/direct access
 - OS reads “size” bytes from file position “from” into “bufAddress”
for (i = from; i < from + size; i++)
bufAddress[i - from] = file[i];
- **Read(fileID, size, bufAddress)** - sequential access
 - OS reads “size” bytes from current file position, fp, into “bufAddress” and increments current file position by size
for (i = 0; i < size; i++)
bufAddress[i] = file[fp + i];
fp += size;



OS File Operations

- **Write** is similar to reads, but copies from the buffer to the file.
- **Seek** just updates fp.
- **Memory mapping** a file
 - Map a part of the portion virtual address space to a file
 - Read/write to that portion of memory \implies OS reads/writes from corresponding location in the file
 - File accesses are greatly simplified (no read/write call are necessary)



File Access Methods

- Common file access patterns from the programmer's perspective
 - **Sequential:** data processed in order, a byte or record at a time.
 - Most programs use this method
 - **Example:** compiler reading a source file.
 - **Direct:** address a block based on a key value.
 - **Example:** database search, hash table, dictionary
- Common file access patterns from the OS perspective:
 - **Sequential:** keep a pointer to the next byte in the file. Update the pointer on each read/write.
 - **Direct:** address any block in the file directly given its offset within the file.



Naming and Directories

- Need a method of getting back to files that are left on disk.
- OS uses numbers for each files
 - Users prefer textual names to refer to files.
 - **Directory:** OS data structure to map names to file descriptors
- Naming strategies
 - **Single-Level Directory:** One name space for the entire disk, every name is unique.
 1. Use a special area of disk to hold the directory.
 2. Directory contains <name, index> pairs.
 3. If one user uses a name, no one else can.
 4. Some early computers used this strategy. Early personal computers also used this strategy because their disks were very small.
 - **Two Level Directory:** each user has a separate directory, but all of each user's files must still have unique names



Naming Strategies (continued)

- Multilevel Directories - tree structured name space (Unix, and all other modern operating systems).
 1. Store directories on disk, just like files except the file descriptor for directories has a special flag bit.
 2. User programs read directories just like any other file, but only special system calls can write directories.
 3. Each directory contains <name, fileDesc> pairs in no particular order. The file referred to by a name may be another directory.
 4. There is one special root directory. *Example:* How do we look up name: `/usr/bin/l`
- Limitations with basic tree structure
 - Difficult to share file across directories and users
 - Can't have multiple file names



Directory Operations

- Search for a file: locate an entry for a file
- Create a file: add a directory listing
- Delete a file: remove directory listing
- List a directory: list all files (*ls* command in UNIX)
- Rename a file
- Traverse the file system



Protection

- The OS must allow users to control sharing of their files => control access to files
- Grant or deny access to file operations depending on protection information
- **Access lists and groups** (Windows NT)
 - Keep an access list for each file with user name and type of access
 - Lists can become large and tedious to maintain
- **Access control bits** (UNIX)
 - Three categories of users (owner, group, world)
 - Three types of access privileges (read, write, execute)
 - Maintain a bit for each combination (111101000 = rwxr-x---




Summary of File System Functionality

- Naming
- Protection
- Persistence
- Fast access



How Disks Work





- The disk surface is circular and is coated with a magnetic material. The disk is always spinning (like a CD).
- Tracks are concentric rings on disk with bits laid out serially on tracks.
- Each track is split into *sectors* or *blocks*, the minimum unit of transfer from the disk.



How Disks Work

- CDs come individually, but disks come organized in *disk pack* consisting of a stack of platters.
- Disk packs use both sides of the platters, except on the ends.
- Comb has 2 read/write head assemblies at the end of each arm.
- *Cylinders* are matching sectors on each surface.
- Disk operations are in terms of radial coordinates.
 - Move arm to correct track, waiting for the disk to rotate under the head.
 - Select and transfer the correct sector as it spins by



Disk Overheads

- **Overhead:** time the CPU takes to start a disk operation
- **Positioning time:** the time to initiate a disk transfer of 1 byte to memory.
 - **Seek time:** time to position the head over the correct cylinder
 - **Rotational time:** the time for the correct sector to rotate under the head
- **Transfer rate:** once a transfer is initiated, the rate of I/O transfer (bandwidth)



File Organization on Disk

The information we need:

fileID 0, Block 0 → Platter 0, cylinder 0, sector 0

fileID 0, Block 1 → Platter 4, cylinder 3, sector 8

...

Key performance issues:

1. We need to support sequential and random access.
2. What is the right data structure in which to maintain file location information?
3. How do we lay out the files on the physical disk?

