Silberschatz and Galvin
Chapter 14

Tertiary Storage Structure

Tertiary Storage Structure

- Tertiary storage devices
- Operating system issues
- Performance issues

Tertiary Storage Devices

- Defining characteristic: low cost
- Generally built using *removable media*
- Examples: floppy disks, CD-ROM, …
  - Floppy disks: thin flexible disk coated with magnetic material, enclosed in a protective plastic case
  - Optical disks: materials that are altered by laser light to have spots that are relatively light and dark
    - Phase-change disk: crystalline or amorphous state
    - Dye-polymer disk: laser heat makes bumps, warms bumps to remove them

Magneto-optic disk

- Magnetic material covered with protective layer of plastic or glass; head much farther from disk than with magnetic disk; less susceptible to head crashes
- Laser heat makes spot susceptible to magnetic field (records)
- Laser light polarization when bouncing off of magnetic spot used for reading (Kerr effect)

Removable disks

- Read-write disks
  - Magnetic disks, magneto-optic disks, optical disks
- Write-once, read many (WORM)
  - One example: thin aluminum film sandwiched between two glass or plastic platters; holes burnt through aluminum; information can be destroyed but not altered
  - Another example: CD-R
- Read-only disks
  - Examples: CD-ROM and DVD

Robotic jukebox for Magneto-optical disks

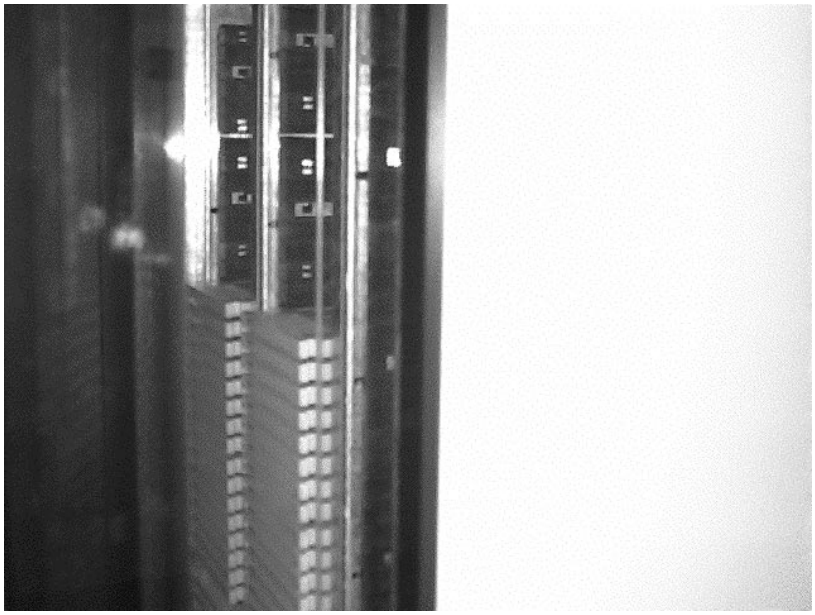

Magnetic tape

- Compared to disk:
  - Less expensive, holds more data, random access much slower
  - Robotic tape installations
    - Stacker: library that holds a few tapes
    - Silo: library that holds thousands of tapes
  - *Archive* disk resident tapes for low-cost storage. *Stage* back into disk storage for active use

Magnetic tape

- Tape silo at Jefferson Lab (http://www.jlab.org/ccc/silo/info/)
  - One terabyte of data a day received
  - 6000 tapes
  - 50 gigabytes per tape at present
  - 300 terabytes total storage
  - Expected enhancements up to over a petabyte (1,000 terabytes) of "near-line" storage

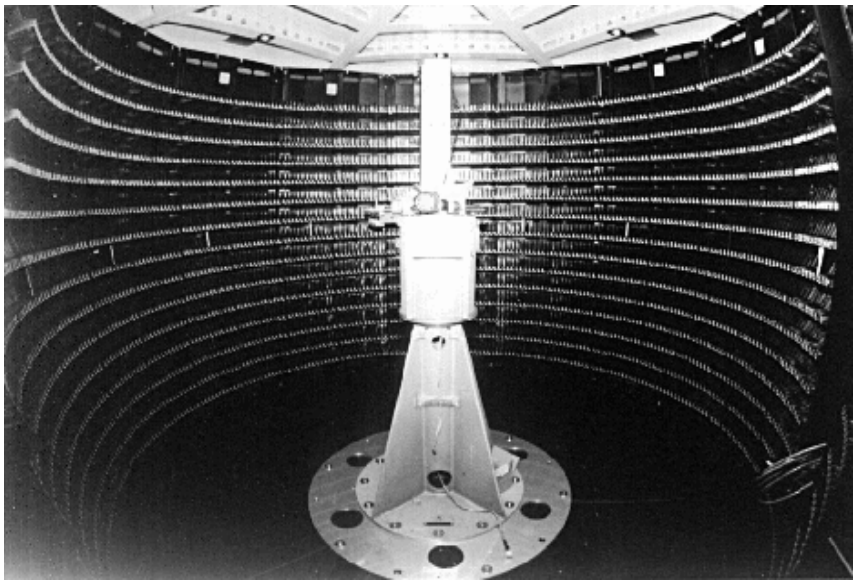

Interior of the STK 4410 before the robotics arm was installed.

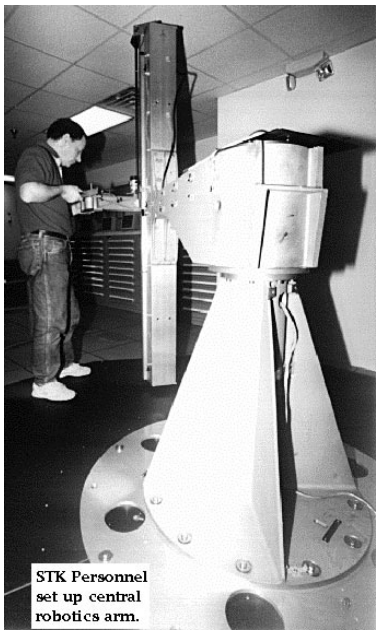

STK Personnel set up central robotics arm.

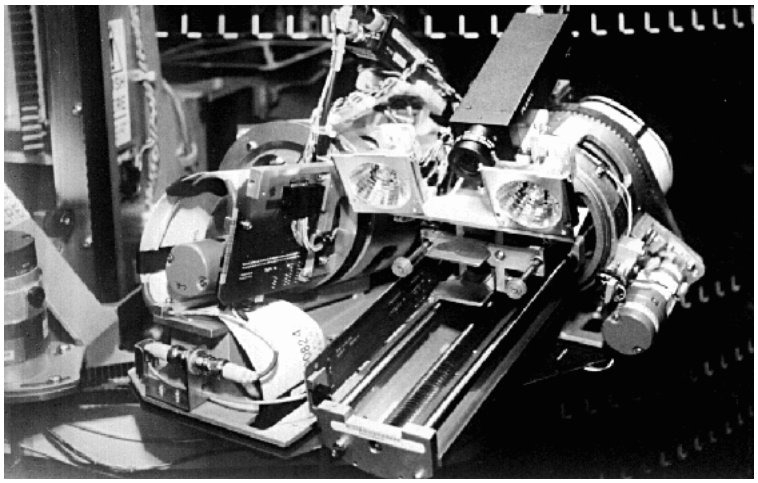

The robot's "head" contains two cameras and two hands for information retrieval.

Operating System Issues

- Major OS jobs
  - Manage physical devices
  - Present a virtual machine abstraction to applications
- Hard disk abstractions
  - Raw device: an array of data blocks
  - File system: OS queues and schedules the interleaved requests from several applications

Application interface

- Most OSs handle removable disks almost exactly like fixed disks--a new cartridge is formatted and an empty file system is generated on the disk.
- Tapes: raw device. Application opens whole tape drive rather than file on tape.
  - Tape drive reserved for exclusive use of application
  - Application decides how to use the array of blocks
  - Tape's format is generally specific to the program that created it

Tape drives

- Basic operations for tape drives differ from those of a disk drive
- **Locate**: position tape to specific logical block (instead of **seek**)
  - Locate 0 is the same as rewinding
- **Read position**: current logical block
- **Space**: relative movement over logical blocks
  - Space -2: go back two logical blocks
- Append-only devices. Update effectively erases everything past that block
- EOT mark follows last block on tape

Speed

- Tertiary storage aspects of speed: *bandwidth* and *latency*
- Bandwidth: measured in bytes per second
  - *Sustained bandwidth*--average data rate during a large transfer; number of bytes/transfer time (this is the data rate when the data stream is actually flowing)
  - *Effective bandwidth*--average over the entire I/O time, including **seek** or **locate**, and cartridge switching (this is the drive's overall data rate)

Speed

- Access latency--amount of time needed to locate data
  - Access time for a disk–move the arm to the selected cylinder and wait for the rotational latency; generally less than 35 milliseconds
  - Access time on tape requires winding tape reels until the selected block reaches the tape head; tens or hundreds of seconds.
  - Generally say that random access within a tape cartridge is about a thousand times slower than random access on disk.
- Access times on jukebox or tape silo (robotic arm) also requires time to remove (including possibly a return to a consistent state), locate, and load media. Hence removable library best for infrequently used data.

Reliability

- Fixed disk drive is likely to be more reliable than removable disk or tape drive
- Optical cartridge is likely to be more reliable than a magnetic disk or tape
- Head crash in a fixed hard disk generally destroys the data whereas the failure of a tape drive or optical disk drive often leaves the data cartridge unharmed
- Recently, much controversy over lifetimes of CD-ROM
  - Manufactured CD-ROMs versus CD-R (predictions in both directions)
  - Years or decades?

Cost

- Main memory is much more expensive than disk storage
- The cost per megabyte of hard disk storage is competitive with magnetic tape if only one tape is used per drive
- The cheapest tape drives and the cheapest disk drives have had about the same storage capacity over the years
- Tertiary storage gives a cost savings only when the number of cartridges is considerably larger than the number of drives