Chapter 10

Computer Storage Fundamentals

Ed Crowley

Spring 10

Intro

- Understanding storage basics requires understanding basic computer:
  - Hardware
  - Software
  - Operating systems

PC Block Diagram

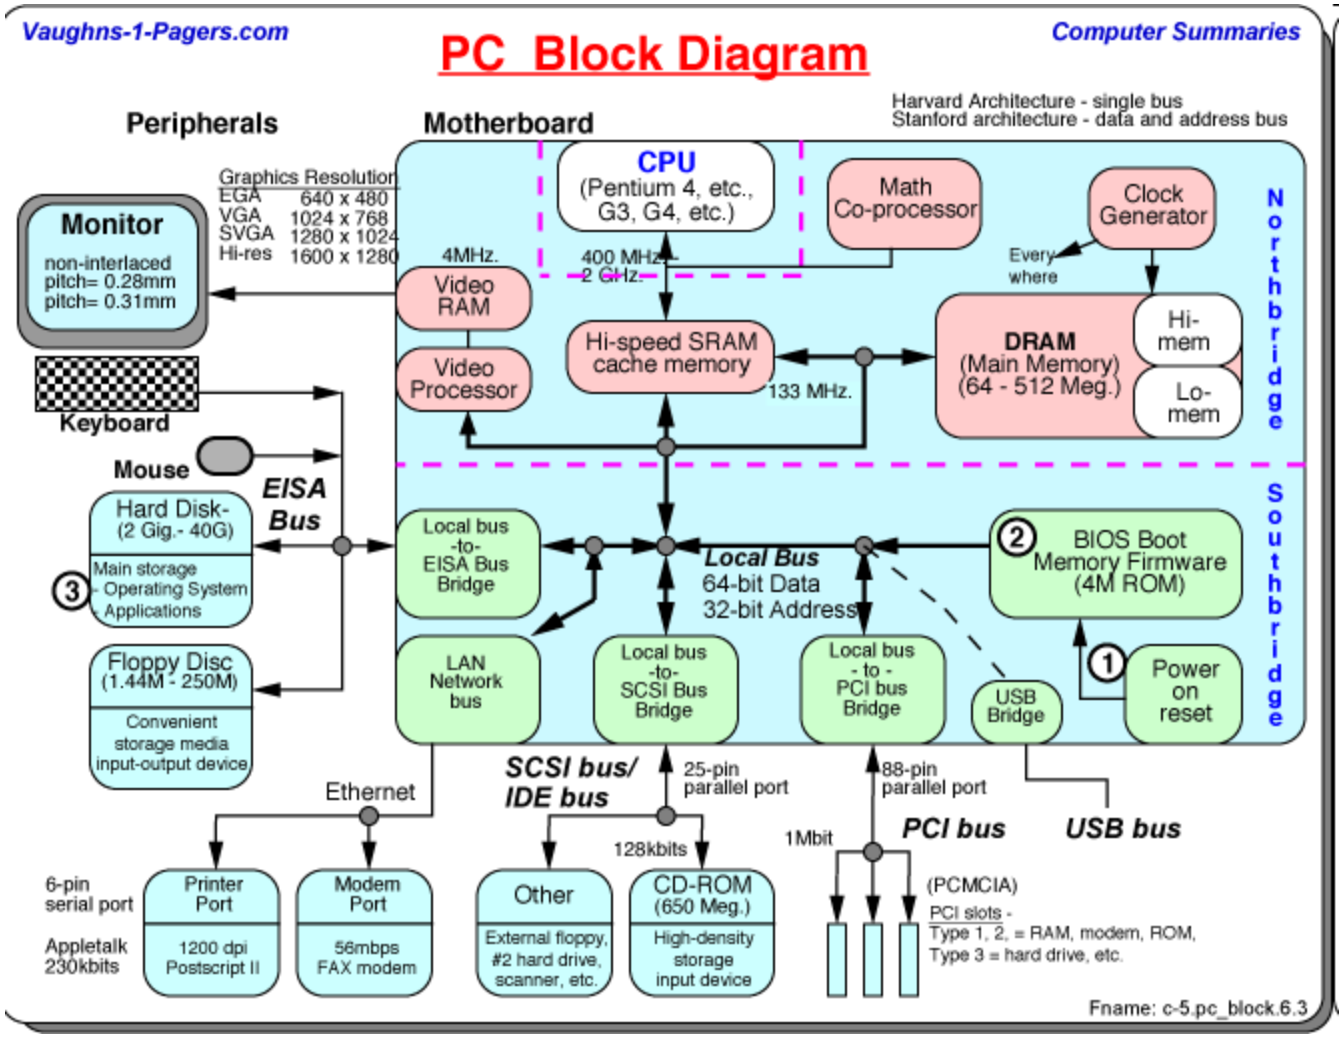

http://www.vaughns-1-pagers.com/computer/pc-block-diagram.htm

Historical Hard Drives and Interfaces

- Two primary drive interface standards
  1. IDE/ATA (Integrated Drive Electronics/AT Attachment) and
  2. SCSI (Small Computer System Interface) are the
  - There are others, but these are our focus.
- ATA-1 was designed with a single data channel that could support two hard drives, one jumpered as master and the other as slave.
  - Originally supported programmed I/O (PIO) modes 0 through 2, with a maximum transfer rate of 8.3 MB/sec
  - Multiple IDE/ATA modes are backwards compatible

Note that this interface is now sometimes referred to as the Parallel ATA Interface (PATA).

Drive Size Boundaries

- Historically, different drive technology generations utilized different technologies with different boundaries
  - Historical boundaries included: 2.1 GB, 8.4 GB, and 32 GB
- Current ATA interface standard uses 28-bit addressing
  - Tops out at approximately 137.4GB.
- Drive utilities do not currently operate properly on drives over 137GB, even if the ATA controller card is updated.
- For more current information:

  `http://www.pcguide.com/ref/hdd/if/ide/std.htm`

Drive Cabling

- A goal is getting the maximum amount of data safely transferred in the least amount of time.

Cable requirements

- Up until the ATA/33 standard, a 40 conductor/40 pin cable was sufficient.
- Anything faster than ATA/33 requires an 80 conductor/40pin cable.
  - Blue connector to the host controller on motherboard
  - Gray connector to slave
  - Black connector to the master

Mixed HD and Cable Types

- You can place ATA drives with different ratings on the same cable.
  - When ATA/33 is combined with ATA/100 everything defaults to the slower speed.
- The entire IDE/ATA system is now designed for cable select mode to operate as initially planned
- ATA bridges make hardware based write protection and the capability to hot swap hard drive media.

Serial ATA (SATA)

- August 2001, released
  - Parallel ATA interface replacement .
- Offers:
  - Backwards compatibility for existing ATA and ATAPI devices
  - Thin small cable solution
    - Makes easier cable routing
  - Better airflow compared to ATA ribbon (parallel) cables.

Serial ATA vs. Parallel Connectors

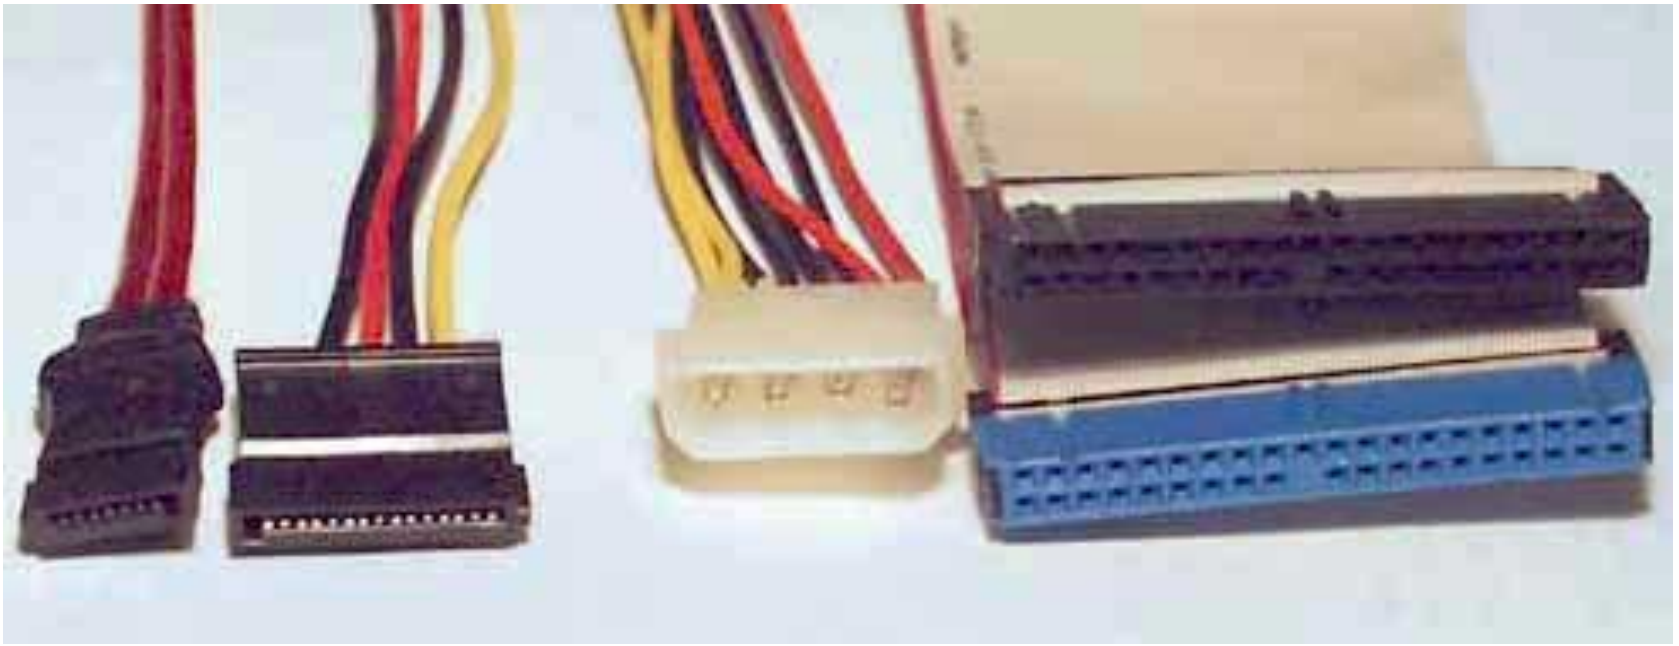

SCSI

- Typically, IDE and SCSI drives use identical head and platter assembly.
  - Different controllers.
- For the duration of a read or write operation, ATA devices utilize the entire bus.
  - At any one time on each ATA bus, only one ATA device can be active.
- If the O/S supports it, SCSI can support parallel, queued commands where multiple devices can be used at once.

SCSI Standards

SCSI Standard	Common Name	External Transfer Speed	Cable Type
SCSI-1	Asynchronous	4 MB/s	A (50 pin)
SCSI-2	Wide	10 MB/s	P (68 pin)
SCSI-2	Fast	10 MB/s	A (50 pin)
SCSI-2	Wide / Fast	20 MB/s	P (68 pin)
SCSI-3	Ultra / Wide	20 / 40 MB/s	P (68 pin)
SCSI-3 pin)	Ultra2 / Wide	40 / 80 MB/s	A or P (50/68
SCSI-3	Ultra3 / Ultra160	160 MB/s	P (68 pin)
SCSI-3	Ultra4 / Ultra320	320 MB/s	P (68 pin)

SCSI Signaling Types

- Single Ended
  - Original, short cable lengths, one signal, one ground
- Low voltage differential
  - Boosted signal, put inverse on second wire to eliminate interference problems
- Low voltage differential/multimode
  - Economical balance between SE and HVD
- High voltage differential
  - LVD/SE devices are compatible with SE or LVD signaling

SCSI Cables and Connectors

Cable A

- 50 conductors and supports standards built on 8-bit transfer widths.

Cable P

- 68 conductors and supports standards built on 16-bit transfer widths.
  - When you mix cable types by using 50-pin to 68-pin adapter plugs, the throughput of the entire chain will drop to the lower rate.
- Any time that signaling cables use a relatively high voltage to transfer information, they are susceptible to signal reflection from the end of the cable.

Preparation of Hard Drive Media

- Prior to forensic duplication, hard drives need to be prepped
- When analysis is performed through virtual disk mounting, the nature of the duplicated image file will keep old data from falling within the scope of your investigation.
- The following use virtual disks to expose the evidence files to the examiner.
  - The Linux Loopback device
  - OnTrack's Forensic Toolkit
  - ASRData's SMART suite
  - Guidance Software's EnCase

Prepping with dd

- Target drive must be prepped
- Clean (prep) storage media with dd.
  - dd command copies blocks of data from its “in file” to its “out file”.
  - In this case, we use the /dev/zero device as the source, because this will give us a continuous source of NULL values (hex char 0x00).

```
# dd if=/dev/zero of=/dev/hdb
```

Partitioning and Formatting Storage Drives

- For the most reliable results when formatting a partition, use an operating system native to the format that you intend to use.
  - With XP, drives are formatted with the Disk Management console.
  - For Linux, you can verify that the hard drive (target drive) has been recognized by the BIOS and the operating system by running the dmesg command.

```
Dmesg | grep hd
Fdisk /dev/hdb
```

Introduction to File Systems and Storage Layers

Six file system layers

- Physical
- Data classification
- Allocation units
- Storage space management
- Information classification
- Application-level storage

Six File System Layers

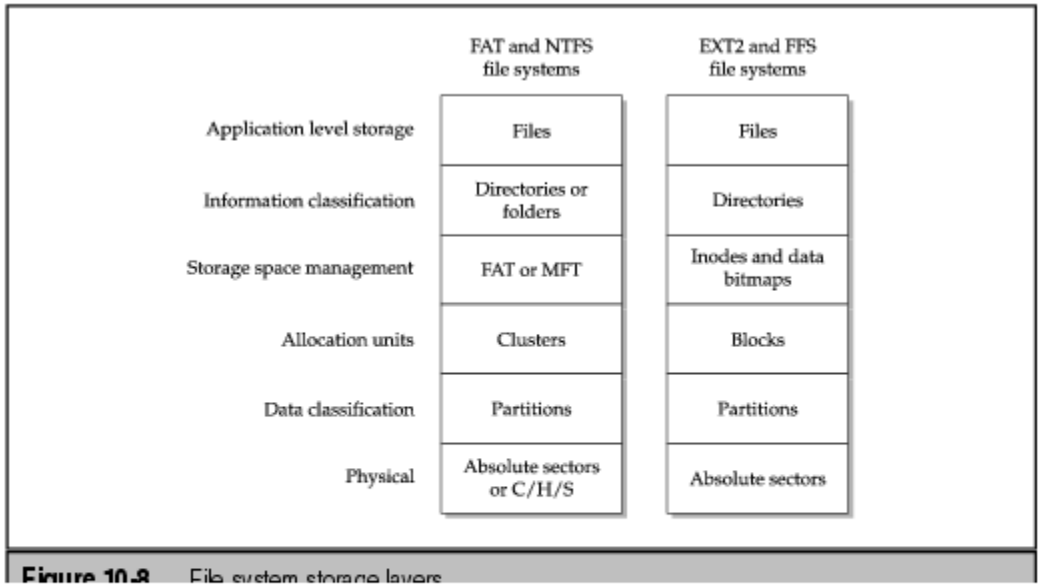

Figure 10-8 File system storage layers

Physical Layer

- Drive hardware reads and writes in 512 byte blocks (sectors).
  - Starting 2011, drive sectors will be 4096 bytes.
- Absolute sectors are numbered sequentially
  - Start at zero and continue until end of drive.
- Intel hardware exposes an additional interface that uses three values (cylinder, head, and sector, CHS) to locate a specific portion of the disk.

Data Classification Layer

- If Windows observes an unknown partition type ID, it ignores the partition.
  - Even if it is formatted correctly and has valid data.
  - Most Unix variants use the term slice rather than partition.
- In order to remain compatible with the Intel partitioning specification, BSD partitions tables and slices are encapsulated within a standard partition scheme.
- This means that the entire file system for a new installation of FreeBSD on an Intel based computer will actually reside within the first partition on the hard drive.

The Allocation Units Layer

- Blocking is the allocation method used by the operating system.
- The size of each allocation unit depends on three variables:
  - Type of file system
  - Partition size
  - Administrator knowledge
- Each file system defines its own scheme for laying out data on the storage medium

Storage Space Management Layer

- Layer manages the thousands of allocation units present on a file system.
  - Allocation unit is the smallest addressable chunk of data that the operation system can handle.
- FAT file systems use a file allocation table (FAT) to track the status of every file system allocation unit.

Information Classification and Application-level Storage Layers

- Top two layers of the file system storage model consist of directories and files.
  - Generally, you can reduce the number of files that require examination by comparing hash values to known good file lists.
- Another method used to minimize spurious data is elimination based on file timestamps.

Questions?