I/O Hardware

COMP755 Advanced Operating Systems

I/O Devices

Disk	Keyboard
hard	Speakers
floppy	Printer
CD	Mouse
DVD	Scanner
Monitor	Game controller

Many, many more

I/O Controller

- Connects the I/O devices to the system.
- Communicates with the CPU and the RAM over the bus.
- A single I/O controller may control multiple devices.
- Most computers have several I/O controllers.
- Actions are initiated by the CPU.

Differences in I/O Devices

- Complexity of control
- Unit of transfer
  - stream of bytes
  - large blocks
- Data representation
  - Encoding schemes
- Error conditions
  - Devices respond to errors differently

Differences in I/O Devices

- Programmed I/O
  - Process is busy-waiting for the operation to complete
- Interrupt-driven I/O
  - I/O command is issued
  - Processor continues executing instructions
  - I/O module sends an interrupt when done

Differences in I/O Devices

- Data rate
  - May be differences of several orders of magnitude between the data transfer rates
  - High speed devices may run faster than the CPU or bus can process the input.

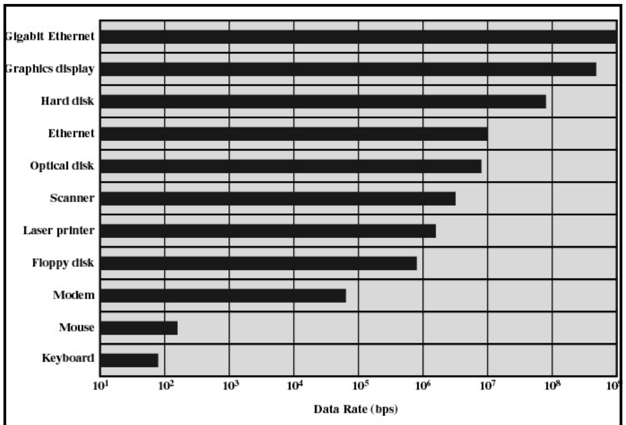

Figure 11.1 Typical I/O Device Data Rates

Techniques for Performing I/O

- Direct Memory Access (DMA)
  - DMA module controls exchange of data between main memory and the I/O device
  - Processor interrupted only after entire block has been transferred
  - DMA may be built into the device controller or may be a separate bus controller

Direct Memory Access

- Takes control of the system from the CPU to transfer data to and from memory over the system bus
- Cycle stealing is used to transfer data on the system bus
- The instruction cycle is suspended so data can be transferred
- The CPU may pause one bus cycle
  - No interrupt occurs
- Impact diminished by CPU caching

Buffer Chaining

- Handles multiple transfers without the processor
- Device given linked list of buffers
- Device hardware uses next buffer on list automatically

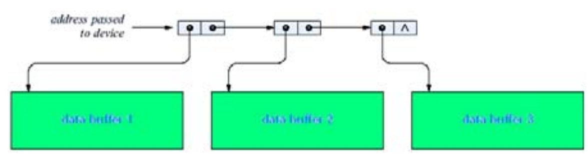

Scatter Read and Gather Write

- Special case of buffer chaining
- Large data transfer formed from separate blocks
- *Example*: to write a network packet, combine packet header from buffer 1 and packet data from buffer 2
- Eliminates application program from copying data into single, large buffer

Operation Chaining

- Further optimization for smart device
- Processor gives series of commands to device, sometimes called a channel program
- Device carries out successive commands automatically

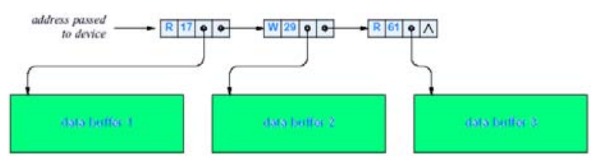

Disk Performance Parameters

Disk read or write involves three factors

1. Seek time
   – time it takes to position the head at the desired track
2. Rotational delay or rotational latency
   – time its takes for the beginning of the sector to reach the head
3. Transfer time
   – time required for the data to move under the head

WD Caviar® | 80 GB 7200 RPM

Rotational Speed	7,200 RPM (nominal)
Average Read Seek Time	8.9 ms
Track-To-Track Seek Time	2.0 ms (average)
Full Stroke Seek	21.0 ms (average)
Average Rotational Latency	4.2 ms (nominal)
Cylinders	16,383
Number of Heads (Physical)	6
Sectors Per Track	63
Bytes Per Sector	512

Performance Example

How long does it take to read a 512 byte block from the disk?

Average Seek time	8.9 ms
Average Rotational Delay	4.2 ms
Transfer time	$1blk * \frac{60\,sec/\min}{7200\,rev/\min * 63\,blk/rev} = 0.13\,m\sec$
Total	13.2 ms

Performance Example

How long does it take to read two consecutive 512 byte blocks from the disk?

Average Seek time	8.9 ms
Average Rotational Delay	4.2 ms
Transfer time	$2blk * \frac{60\sec/\min}{7200\,rev/\min * 63\,blk/rev} = 0.26\,m\sec$
Total	13.4 ms

Logical vs. Physical

- Many disks present the OS with a logical layout that is different from the physical layout.
- Compatibility with old standards
- Many disks use Logical Block Addressing (LBA) to hide the physical layout.

CD Performance

- The “X” of CD speed claims represents the number of times faster the CD spins than music CD players.
- CDs are written in a long spiral instead of concentric tracks.
- The average seek time for a CD is 90 ms (compare to 9 ms for a hard drive).
- CDs are efficient reading large files but slow reading small files.

RAID

- Redundant Array of Independent Disks
- A collection of disks are used as one large unit of mass storage.
- Multiple disks operating simultaneously can increase the data transfer rate.
- Extra data stored on the disks can recover the information should a disk fail.

RAID Types

- RAID 0 - Striping
- RAID 1 - Mirroring
- RAID 2 - Hamming code error recovery
- RAID 3 - Bit-interleaved parity
- RAID 4 - Block-level parity
- RAID 5 - block-level distributed parity
- RAID 6 - Dual redundancy

RAID 0 (non-redundant)

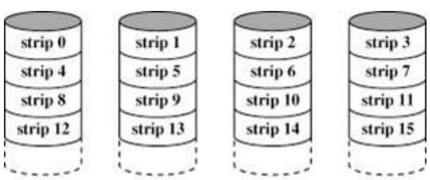

(a) RAID 0 (non-redundant)

- •Improved transfer rate
- •Decreased reliability

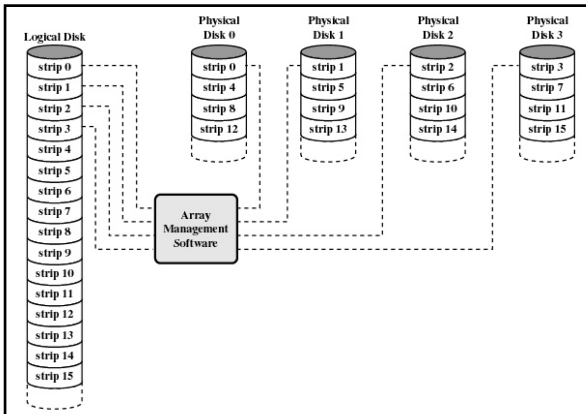

Figure 11.10 Data Mapping for a RAID Level 0 Array [MASS97]

RAID 1 (mirrored)

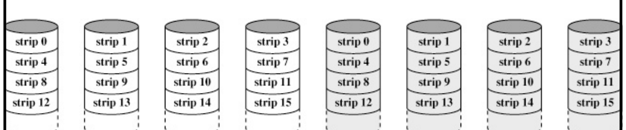

b) RAID 1 (mirrored)

Figure 11.9 RAID Levels (page 1 of 2)

- •Improved Reliability
- •Slightly slower writes.
- •Possibly faster reads

RAID 2 (redundancy through Hamming code)

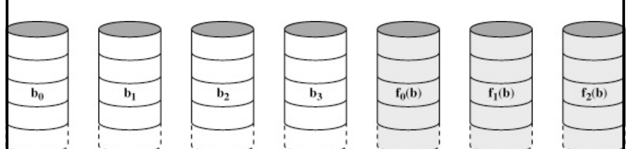

c) RAID 2 (redundancy through Hamming code)

Figure 11.9 RAID Levels (page 1 of 2)

RAID 3 (bit-interleaved parity)

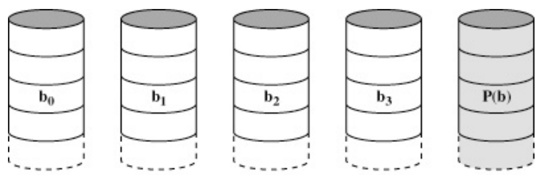

(d) RAID 3 (bit-interleaved parity)

Figure 11.9 RAID Levels (page 2 of 2)

RAID 4 (block-level parity)

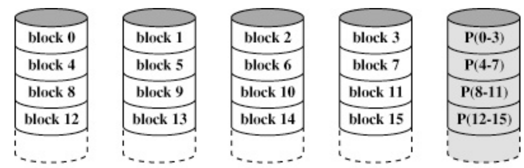

(e) RAID 4 (block-level parity)

Figure 11.9 RAID Levels (page 2 of 2)

RAID 5 (distributed block parity)

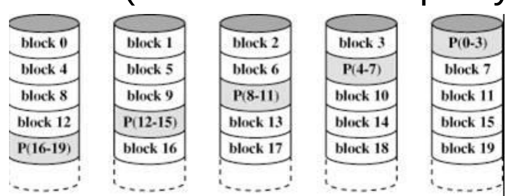

RAID 5 (block-level distributed parity)

- Striping improves read performance
- Parity improves reliability

RAID 6 (dual redundancy)

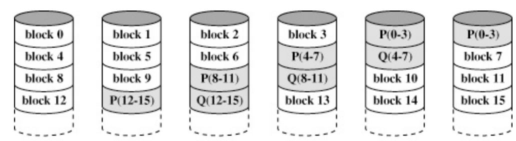

(g) RAID 6 (dual redundancy)

Figure 11.9 RAID Levels (page 2 of 2)

- Like RAID 5 but with two parity blocks for each data block
- Slow writes