Storage and I/O Devices

What is the point behind building fast CPUs, if the speed of I/O does not improve?

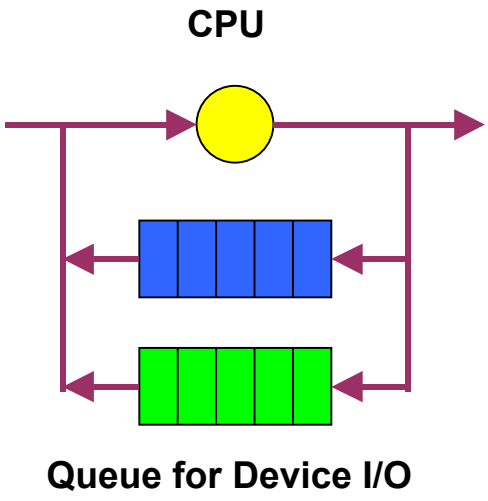

Life of a process

Example

Per **Amdahl Law**, regardless of the CPU speedup, the overall speedup never exceeds 100%!

Disk

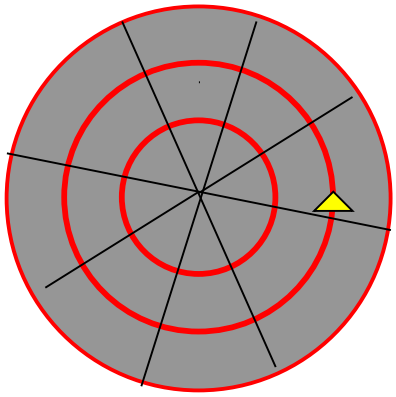

Recognize

Tracks, sectors,

Platter/surfaces

Cylinders etc.

Time to read or write from a disk =

Seek time +

Rotational latency +

Controller time +

Queuing delay +

Data transfer time

Disk Characteristics in 2001

	Seagate Cheetah Ultra160 SCSI	IBM Travelstar	IBM 1 GB Microdrive
Diameter	3.5”	2.5”	1.0”
Formatted capacity	73.4 GB	32.0 GB	1.0 GB
Cylinders	14,100	21,664	7,167
Disks	12	4	1
Heads	24	8	2
Bytes / sector	512-4096	512	512
Sectors / track	~424	~512	~512
RPM	10,033	5,411	3,600
Seek time *	6.0 ms	12.0 ms	12.0 ms
Data transfer rate	27-40 MB/s	11-21 MB/s	2.6-4.2 MB/s

Continued advance in capacity of 60% per year, and bandwidth of 40% per year.

Disk Array: RAID

Disk access time and availability can be improved using disk cache and RAID (Redundant Array of Inexpensive Disks). Available at various levels, 0-6.

RAID 0: No redundancy

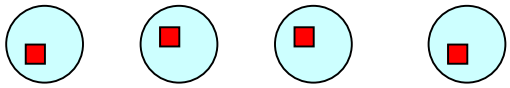

Striping is used to distribute data over several disks. The access time improves due to the possibility of parallel access.

RAID 1: Mirroring

There is 100% redundancy. If one crashes, then data is available from the other

RAID 3: Bit-interleaved parity

Each bit of an N-bit word is written on a separate disk. There is one extra disk (disk N+1), called the parity disk.

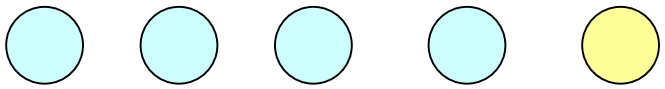

	D0	D1	D2	D3	Parity(even)
Word 0	0	1	1	1	1
Word 1	0	1	0	1	0
Word 2	1	0	1	0	0

If one disk crashes, then data can be reconstructed from the others.

RAID 5: Block Interleaved Distributed Parity

Unlike RAID 4, where all the parity blocks are stored on a separate parity disk, in RAID 5, the parity blocks are distributed. The CRC bits appended to each block handle small read errors

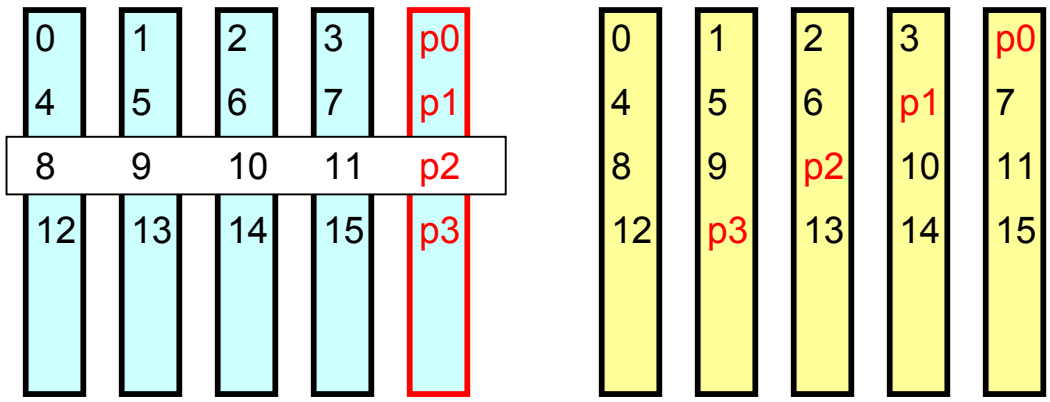

RAID 4 **RAID 5**

Both RAID 4 and RAID 5 allow simultaneous reads, but RAID 5 leads to better parallelism during write operations.

Intelligent Controllers

Intelligent controllers use wider buses and adequate set of buffers to relax the timing constraints and improve efficiency. For example SCSI can support 16 devices on a single bus.

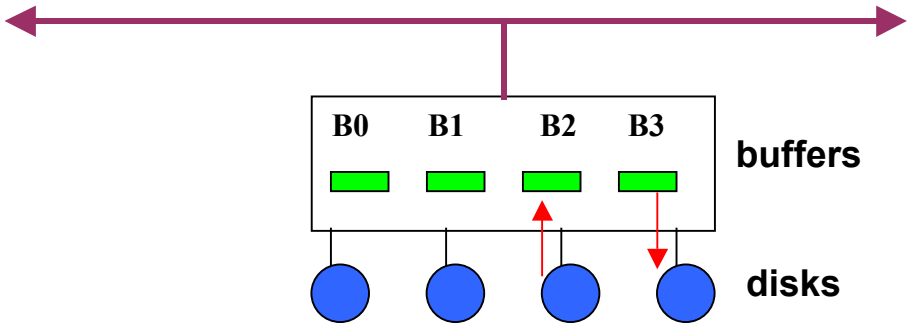

The smartness is due to its ability to handle

- Overlapped commands
- Re-order commands (command queuing)
- Scatter-gather, which provides multiple host addresses in one command.

Example of overlapped command processing

Consider the following queue of requests:

Read Disk 0
Read Disk 1
Write Disk 2
Write Disk 3