

Computer System



2

Why I/O Matters?


- CPU performance increase $\sim 60\%/Year$
- I/O performance increase $< 10\%/Year$
 - Limited by *mechanical* delays
- *Amdahl's Law*: system speed-up limited by the slowest part!
- Example:
 - Suppose 1 sec I/O + 4 sec CPU \Rightarrow 5 seconds
 - Increasing CPU by 100% \Rightarrow 3 seconds
 - We only see 66% Speedup \Rightarrow *I/O bottleneck!*

• "I think Silicon Valley was misnamed. If you look back at the dollars shipped in products in the last decade, there has been more revenue from magnetic disks than from silicon. They ought to rename the place Iron Oxide Valley."

--Al Hoagland, 1982

4

Typical Collection of I/O devices



- Bus is the connection between Processor, Memory, and I/O devices and processor
- Protocols on the bus and interrupts used for communication among devices and processor

Lecture 21: Input/Output

COS 471a, COS 471b / ELE 375

Computer Architecture and Organization

Princeton University

Fall 2004

Mario Ibrahim for Prof. David August

1

3

Types & Characteristics of I/O Devices

- Behavior
 - Input *read once*
 - Output *write only, cannot be read*
 - Storage *can be reread and usually rewritten*
- Partner
 - What's on the other end? *Human or Machine*
- Data Rate
 - Peak Rate of transfer *between I/O and Memory or CPU*
- Keyboard -> Input Device -> Used by Human -> 10 B/s

6

Measuring I/O Performance

- Primary measures of performance:
 - Bandwidth – *throughput*
 - How much data can we move through the system in a certain time?
 - Used to measure performance in supercomputer applications, large streams of data
 - How many I/O operations can we do per unit of time?
 - Application for tax processing office, process large number of small forms
 - Units
 - Bandwidth, *base 10*
 - 1 KByte/sec = 1,000 Bytes/sec
 - 1 MByte/sec = 1,000,000 Bytes/sec
 - Memory, *base 2*
 - 1 KByte = 2^{10} = 1024 Bytes
 - 1 MByte = ?
 - 1 GByte = ?

5

Mouse

- Why "mouse"?
 - "It would be wonderful if I can inspire others, who are struggling to realize their dreams, to say 'if this country klu could do it, let me keep slogging away'." - *Douglas Engelbart*
- Use *optical* or *mechanical* means to determine movement
 - Generates pulses when moved (using LED & detector in Ch. 1)
 - Updates counters
 - Why does motion appear smoothly?
 - We are slow relative to processor rate of reading mouse status
- Bandwidth requirement limited by human hand coordination




8

Diverse I/O Devices

Device	Behavior	Partner	Data Rate (KB/sec)
Keyboard	input	human	0.01
Mouse	input	human	0.02
Monitor	output	human	60,000
Laser Printer	output	human	200
Modem	input or output	machine	2 - 8
Network/LAN	input or output	machine	500 - 6,000
Optical Disk	storage	machine	1,000
Magnetic Disk	storage	machine	2,000 - 10,000

7

Organization of Magnetic Disk



- Typical Numbers
 - 1 to 15 Platters
 - two recordable sides each
 - 3600 to 7200 RPM
 - 1 to 8 inches in diameter
 - 1,000 to 5,000 tracks per surface
 - 64 to 200 sectors per track



10

Magnetic Disks




- Purpose
 - Non-volatile storage
 - Lowest level of memory hierarchy – *large, cheap, slow*
- Basic Idea
 - Use a rotating platter coated with a magnetic surface
 - *Movable read/write head* accesses the disk
- How does it access data?
 - OS must direct the disk through 3-stage process
 - *Seek time (ST)*– position the head over the proper track
 - *Rotational Latency/Delay (RT)*– desired sector under R/W heads
 - *Transfer time (TT)*– transfer block of bits
 - *Disk controller* manages disk to memory transfer

Disk Access = seek time + rotational delay + transfer time + controller overhead

9


Networks

- *Medium* to communicate between computers
- Characteristics
 - Distance: *0.01 to 10,000 km*
 - Speed: *0.001 to 100 MB/sec*
 - Topology: *Bus, Ring, Star, Tree*
- Examples
 - RS232 standard – star topology, slow
 - LAN – bus topology, 10 Mbit/sec



12

A Closer Look



11

Difficulties in Designing a Bus

- Physical restraints that limit the maximum bus speed:
 - Length of Bus (wires)
 - Number of devices attached to the bus
- Various techniques to increase performance
 - *May* adversely affect response time
 - Buffering increases bus bandwidth
 - ...but also increases delay to complete bus access!
- Devices attached to the bus have very different latencies and data transfer rates.

14

Communication Schemes

- Synchronous buses
 - Clock in the control lines and a fixed protocol for communication
 - Relatively fast because protocol is predetermined
 - Disadvantages
 - All bus devices run on same clock rate
 - Clock skew problem constraint on length of bus (Appendix B)
- Asynchronous
 - Have no clock signal
 - Wide variety of devices running at different speeds
 - Uses handshaking protocol
 - Additional control lines needed: *ReadReq*, *DataRdy*, and *Ack*

16

Buses

- *Communication link* shared by memory, the processor and I/O devices
- Consists of *control lines* and *data lines*
 - Control lines signal requests and acknowledgements
 - Data lines carry information between source and destination
- Bus Transactions
 - Sending the *address*
 - Receiving or sending *data*
- Advantages
 - *Versatility* – single connection scheme for easy add-ons
 - *Low Cost* – single set of wires shared in multiple ways
- Disadvantages
 - *Communication bottleneck*, bandwidth limits the maximum I/O throughput.
 - Devices will not always be able to use the bus when they need to.

13

Bus Types

- *Processor-memory* Buses
 - Short, high speed
 - Designed to maximize memory-processor bandwidth
- *I/O* Buses
 - Long, support many types of devices
 - Wide range in bandwidth
 - Does not interface directly with memory,
 - Uses processor-memory bus or backplane bus
- *Backplane* Buses
 - Allow processors, memory and I/O devices to coexist on single bus
 - Balance demands of processor-memory communication with demands of I/O device-memory communication
 - Interconnects the circuit boards containing processor, memory and I/O interfaces

15

Obtaining Access to Bus

- *Bus Master* - Processor
 - Controls access to bus
 - Must initiate and control all bus requests
- *Slave* - Memory
 - Responds to read and write requests
- Drawback of using single master
 - Processor involved in all requests
- Alternative Scheme
 - Multiple Bus Masters
 - Mechanism for arbitrating access to the bus needed

18

Bus Standards

Characteristics	PCI	SCSI
Bus type	backplane	I/O
Data Bus Width	32 - 64	8 - 32
Address/Data Multiplexed		
Bus Masters	multiple	multiple
Arbitration	centralized, parallel arbitration	self-selection
Clocking	synchronous (33 - 66 MHz)	async or sync (5 - 10 MHz)
Peak Bandwidth	80 MB/sec	1.5 - 40 MB/sec
Max Devices	1024	7 - 31
Bus Length	0.5 m	2.5 m

20

Increasing Bus Bandwidth

- Increase data bus width
 - Simultaneous transfer of multiple words
 - More lines means more cost
- Separate address and data lines
 - Allows address and data to be sent at the same time
 - Increases complexity
- Block transfers
 - Send a block of data words in quick succession
 - Increases response time for other devices waiting to use bus

17

Bus Arbitration

- Bus Arbiter decides which bus master gets bus next
- Device signals a *Bus Request*
- Waits for *Grant* signal when bus is available
- Device *Release* bus back to arbiter
- Four Classes of Bus Arbitration schemes:
 - Daisy Chain Arbitration
 - Centralized, Parallel Arbitration
 - Distributed Arbitration by Self-Selection
 - Distributed arbitration by collision detection

19