CS61C - Machine Structures

Lecture 15 - Networks

October 18, 2000

David Patterson

http://www-inst.eecs.berkeley.edu/~cs61c/

Review

- Kernel Mode v. User Mode: OS can provide security and fairness
- Syscall: provides a way for a programmer to avoid having to know details of each I/O device
- To be acceptable, interrupt handler must:
  - service all interrupts (no drops)
  - service by priority
  - make all users believe that no interrupt has occurred

Today's Outline

- Review
- Buses
- Why Networks?
- A Simple Example: Derive Network Basics
- Administrivia
- Protocol, Ethernet
- Internetworking, Protocol Suites, TCP/IP
- Conclusion

Recall : 5 components of any Computer

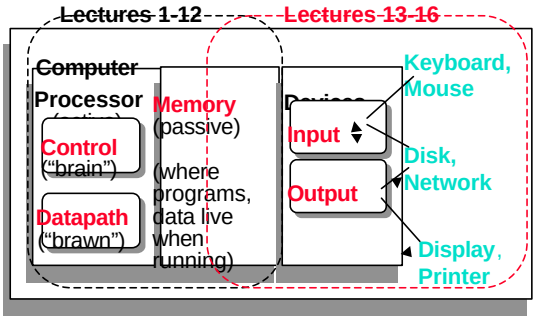

Connecting to Networks (and Other I/O)

- Bus - shared medium of communication that can connect to many devices
- Hierarchy of Buses in a PC

Buses in a PC: connect a few devices

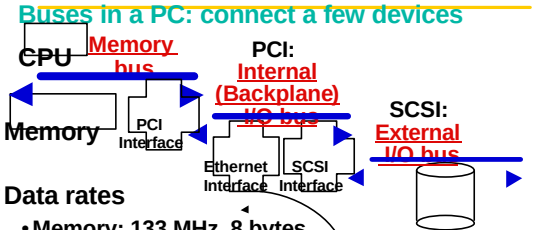

- Data rates
  - Memory: 133 MHz, 8 bytes
    - Þ 1064 MB/s (peak)
  - PCI: 33 MHz, 8 bytes wide
    - Þ 264 MB/s (peak)
  - SCSI: "Ultra3" (80 MHz), "Wide" (2 bytes)
    - Þ 160 MB/s (peak)

Ethernet: 12.5 MB/s (peak)

Why Networks?

- Originally sharing I/O devices between computers
  (e.g., printers)
- Then Communicating between computers
  (e.g, file transfer protocol)
- Then Communicating between people
  (e.g., email)
- Then Communicating between networks of computers
  - Internet, WWW

How Big is the Network (1999)?

~30	Computers in 273 Soda
~400	in inst.cs.berkeley.edu
~4,000	in eecs&cs .berkeley.edu
~50,000	in berkeley.edu
~5,000,000	in .edu
~46,000,000	in US (.com .net .edu .mil .us .org .us)
~56,000,000	in the world

Source: Internet Software Consortium

Growth Rate

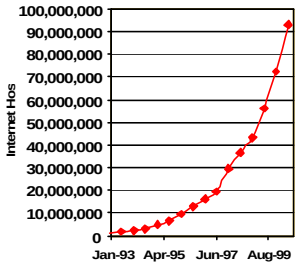

"Source: Internet Software Consortium (http://www.isc.org/)".

Ethernet Bandwidth	
1983	3 mb/s
1990	10 mb/s
1997	100 mb/s
1999	1000 mb/s

What makes networks work?

- links connecting switches to each other and to computers or devices

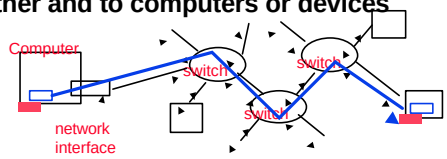

- ability to name the components and to route packets of information - messages - from a source to a destination
- Layering, protocols, and encapsulation as means of abstraction

Typical Types of Networks

- Local Area Network (Ethernet)
  - Inside a building: Up to 1 km
  - (peak) Data Rate: 10 Mbits/sec, 100 Mbits /sec,1000 Mbits/sec (1.25, 12.5, 125 MBytes/s)
  - Run, installed by network administrators
- Wide Area Network
  - Across a continent (10km to 10000 km)
  - (peak) Data Rate: 1.5 Mbits/sec to 2500 Mbits/sec
  - Run, installed by telephone companies
- Wireless Networks, ...

Network Basics: links

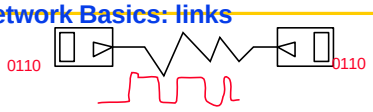

- Link made of some physical media
  - wire, fiber, air
- with a transmitter (tx) on one end
  - converts digital symbols to analog signals and drives them down the link
- and a receiver (rx) on the other
  - captures analog signals and converts them back to digital signals
- tx+rx called a transceiver

Example: Network Media

Twisted Pair: Copper, 1mm think, twisted to avoid antenna effect

Coaxial Cable: Used by cable companies: high BW, good noise immunity

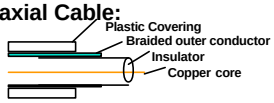

Fiber Optics: Light: 3 parts are cable, light source, light detector

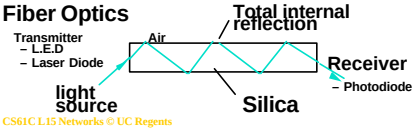

ABCs of Networks: 2 Computers

° Starting Point: Send bits between 2 computers

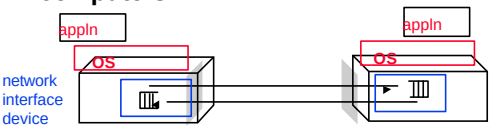

° Queue (First In First Out) on each end

° Can send both ways ("Full Duplex")

° Information sent called a "message"

• Note: Messages also called packets

Administrivia

° Midterm will be Wed Oct 25 5-8 P.M.

• 1 Pimintel

• Midterm conflicts? Talk to TA about taking early midterm ("beta tester")

• 2 sides of paper with handwritten notes; no calculators

• Sample midterm, old midterms online

° Midterm Review Sunday Oct 22 starting 2 PM in155 Dwinelle

° Rest of homework assignments are online: 6, 7, 8

"What's This Stuff Good For?"

In 1974 Vint Cerf co-wrote TCP/IP, the language that allows computers to communicate with one another. His wife of 35 years (Sigrid), hearing-impaired since childhood, began using the Internet in the early 1990s to research cochlear implants, electronic devices that work with the ear's own physiology to enable hearing. Unlike hearing aids, which amplify all sounds equally, cochlear implants allow users to clearly distinguish voices--even to converse on the phone. Thanks in part to information she gleaned from a chat room called "Beyond Hearing," Sigrid decided to go ahead with the implants in 1996. The moment she came out of the operation, she immediately called home from the doctor's office--a phone conversation that Vint still relates with tears in his eyes. ***One Digital Day,*** **1998 (www.intel.com/onedigitalday)**

A Simple Example: 2 Computers

° What is Message Format?

• Similar idea to Instruction Format

• Fixed size? Number bits?

8 bit 32 x Length bits

• Header(Trailer): information to deliver message

• Payload: data in message

• What can be in the data?

• anything that you can represent as bits

• values, chars, commands, addresses...

Questions About Simple Example

° What if more than 2 computers want to communicate?

• Need computer "address field" in packet to know which computer should receive it (destination), and to which computer it came from for reply (source)

8 bits 8 bits 8 bits 32xn bits

Header Payload

ABCs: many computers

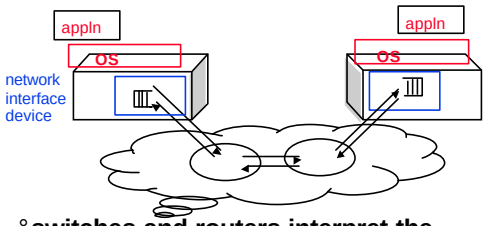

- switches and routers interpret the header in order to deliver the packet
- source encodes and destination decodes content of the payload

Questions About Simple Example

- What if message is garbled in transit?
- Add redundant information that is checked when message arrives to be sure it is OK
- 8-bit sum of other bytes: called "Check sum"; upon arrival compare check sum to sum of rest of information in message

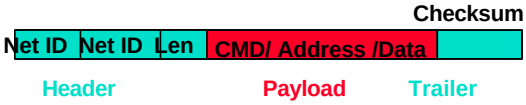

Questions About Simple Example

- What if message never arrives?
- Receiver tells sender when it arrives (ack), sender retries if waits too long
- Don't discard message until get "ACK"; (Also, if check sum fails, don't send ACK)

Observations About Simple Example

- Simple questions such as those above lead to more complex procedures to send/receive message and more complex message formats
- Protocol: algorithm for properly sending and receiving messages (packets)

Ethernet Packet Format

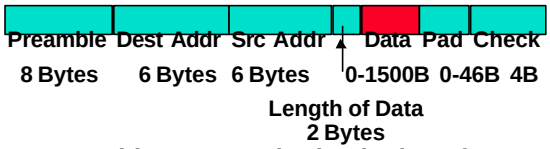

- Preamble to recognize beginning of packet
- Unique Address per Ethernet Network Interface Card so can just plug in & use
- Pad ensures minimum packet is 64 bytes
  - Easier to find packet on the wire
- Header+ Trailer: 24B + Pad

Shared vs. Switched Based Networks

- Shared Media vs. Switched: pairs communicate at same time: "point-to-point" connections
- Aggregate BW in switched network is many times shared
  - point-to-point faster since no arbitration, simpler interface

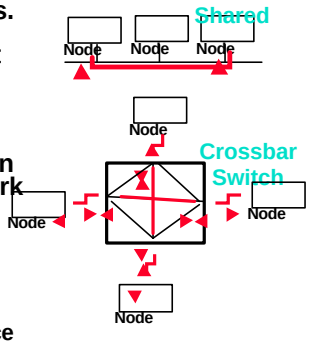

Software Protocol to Send and Receive

- SW Send steps
  - 1: Application copies data to OS buffer
  - 2: OS calculates checksum, starts timer
  - 3: OS sends data to network interface HW and says start
- SW Receive steps
  - 3: OS copies data from network interface HW to OS buffer
  - 2: OS calculates checksum, if OK, send ACK; if not, delete message (sender resends when timer expires)
  - 1: If OK, OS copies data to user address space, & signals application to continue

Protocol for Networks of Networks?

- Internetworking: allows computers on independent and incompatible networks to communicate reliably and efficiently;
  - Enabling technologies: SW standards that allow reliable communications without reliable networks
  - Hierarchy of SW layers, giving each layer responsibility for portion of overall communications task, called protocol families or protocol suites
- Abstraction to cope with complexity of communication vs. Abstraction for complexity of computation

Protocol for Network of Networks

- Transmission Control Protocol/Internet Protocol (TCP/IP)
  - This protocol family is the basis of the Internet, a WAN protocol
  - IP makes best effort to deliver
  - TCP guarantees delivery
  - TCP/IP so popular it is used even when communicating locally: even across homogeneous LAN

Protocol Family Concept

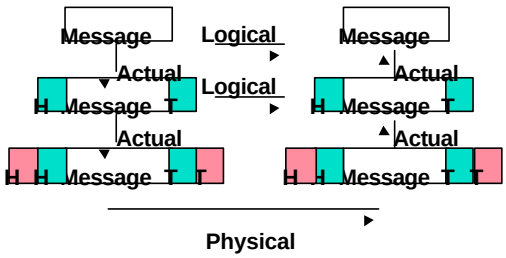

Protocol Family Concept

- Key to protocol families is that communication occurs logically at the same level of the protocol, called peer-to-peer,
- but is implemented via services at the next lower level
- Encapsulation: carry higher level information within lower level "envelope"
- Fragmentation: break packet into multiple smaller packets and reassemble

TCP/IP packet, Ethernet packet, protocols

- Application sends message
- TCP breaks into 64KB segments, adds 20B header
- IP adds 20B header, sends to network
- If Ethernet, broken into 1500B packets with headers, trailers (24B)
- All Headers, trailers have length field, destination, ...

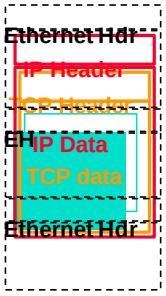

Routing in the Internet

- Individual networks can have own protocols for routing and transmission
- Internet = network of networks
- Designated nodes called gateways know how to route "up" to the backbone based on 'destination nertwork"
- Core gateways know how to route anywhere in the core.

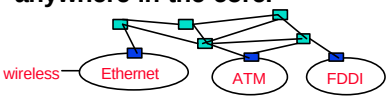

FTP From Stanford to Berkeley (1996)

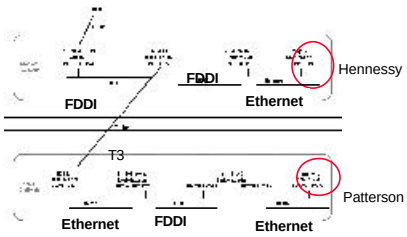

- BARRNet is WAN for Bay Area
  - T3 is 45 Mbit/s leased line (WAN); FDDI is 100 Mbit/s LAN
- IP sets up connection, TCP sends file

What to Remember

- Protocol suites allow heterogeneous networking
  - Another form of principle of abstraction
  - Protocols Þ operation in presence of failures
  - Standardization key for LAN, WAN
- Integrated circuit revolutionizing network switches as well as processors
  - Switch just a specialized computer
- Trend from shared to switched networks to get faster links and scalable bandwidth