

15-213

"The course that gives CMU its Zip!"

Internetworking November 9, 2007

Topics


- Client-server programming model
- Networks
- Internetworks
- Global IP Internet
 - IP addresses
 - Domain names
 - Connectors

lec09-22.ppt

A Client-Server Transaction

Most network applications are based on the client-server model:

- A **server** process and one or more **client** processes
- Server manages some **resource**.
- Server provides **service** by manipulating resource for clients
- Server activated by request from client (vending machine analogy)



Note: clients and servers are processes running on hosts (can be the same or different hosts).

- 2 -

15-213, F07

Computer Networks

A network is a hierarchical system of boxes and wires organized by geographical proximity

- SAN (System Area Network) spans cluster or machine room
Switched Ethernet, Quadrics QSW, ...
- LAN (local area network) spans a building or campus.
Ethernet is most prominent example.
- WAN (wide-area network) spans country or world.
Typically high-speed point-to-point phone lines.

An **internetwork** (**internets**) is an interconnected set of networks.

- The Global IP Internet (uppercase "II") is the most famous example of an internet (lowercase "ii")

Let's see how we would build an internet from the ground up.

- 3 -

15-213, F07

Lowest Level: Ethernet Segment

Ethernet segment consists of a collection of **hosts** connected by wires (twisted pairs) to a **hub**.

Spans room or floor in a building.



Operation

- Each Ethernet adapter has a unique 48-bit address.
- Hosts send bits to any other host in chunks called **frames**.
- Hub blindly copies each bit from each port to every other port. Every host sees every bit.

Note: Hubs are on their way out. Bridges (switches, routers) became cheap enough to replace them (means no more broadcasting).

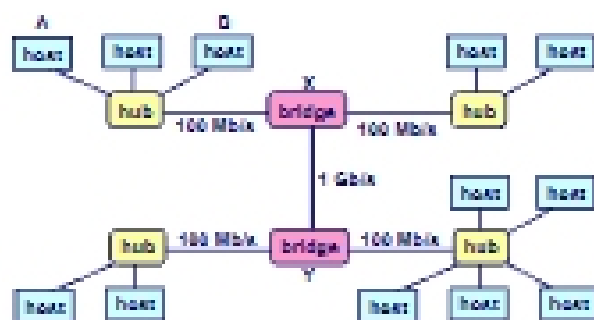
- 4 -

15-213, F07

Next Level: Bridged Ethernet Segment

Spans building or campus.

Bridges cleverly learn which hosts are reachable from which ports and then selectively copy frames from port to port.



- 5 -

15-213, F07

Conceptual View of LANs

For simplicity, hubs, bridges, and wires are often shown as a collection of hosts attached to a single wire:



- 6 -

15-213, F07

Next Level: internets

Multiple incompatible LANs can be physically connected by specialized computers called **routers**.
The connected networks are called an **internet**.




LAN 1 and LAN 2 might be completely different, totally incompatible LANs (e.g., Ethernet and Wifi, 802.11*, 11-linka, DSL, ...)

- 7 -

15-213, F07

Logical structure of an internet



- Ad hoc interconnection of networks
 - No particular topology
 - Vastly different router & link capacities
- Send packets from source to destination by hopping through networks
 - Router forms bridge from one network to another
 - Different packets may take different routes

- 8 -

15-213, F07

The Notion of an internet Protocol

How is it possible to send bits across incompatible LANs and WANs?

Solution: protocol software running on each host and router smooths out the differences between the different networks.

implements an internet protocol (i.e., set of rules) that governs how hosts and routers should cooperate when they transfer data from network to network.

- TCP/IP is the protocol for the global IP Internet.

- 9 -

15-213, F07

What Does an internet Protocol Do?

1. Provides a naming scheme

- An internet protocol defines a uniform format for **host addresses**.
- Each host (and router) is assigned at least one of these internet addresses that uniquely identifies it.

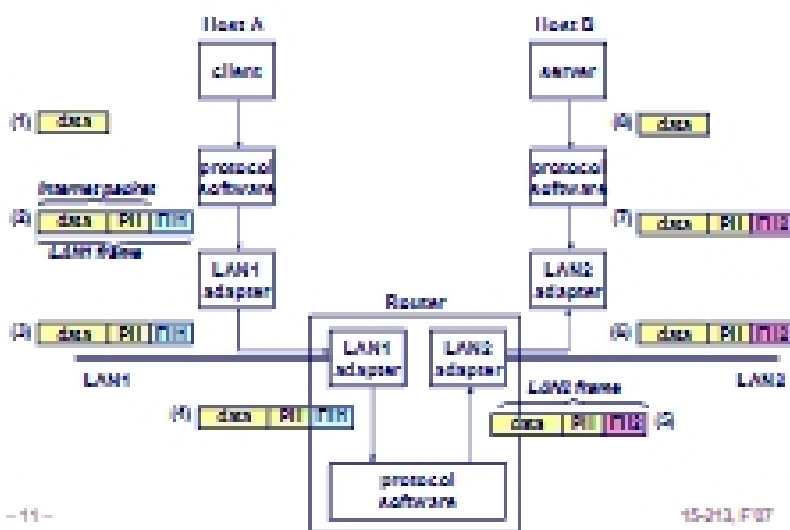
2. Provides a delivery mechanism

- An internet protocol defines a standard transfer unit (**packet**)
- Packet consists of **header** and **payload**
 - Header: contains info such as packet size, source and destination addresses.
 - Payload: contains data bits sent from source host.

- 10 -

15-213, F07


Transferring Data Over an internet



- 11 -

15-213, F07

MCI/WorldCom/UUNET Global Backbone



Source: Boardwatch.com

- 12 -

15-213, F07

Other Issues

We are glossing over a number of important questions:

- What if different networks have different maximum frame sizes? (segmentation)
- How do routers know where to forward frames?
- How are routers informed when the network topology changes?
- What if packets get lost?

These (and other) questions are addressed by the area of systems known as computer networking.

- 13 -

15-213, F07

Global IP Internet

Most famous example of an Internet.

Based on the TCP/IP protocol family

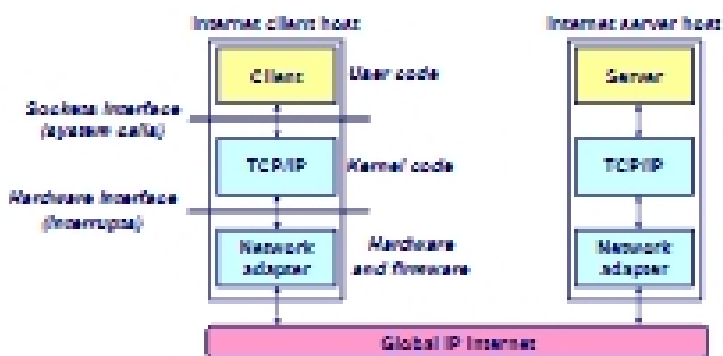
- IP (Internet protocol) :
 - Provides basic naming scheme and unreliable delivery capability of packets (datagrams) from host-to-host.
- UDP (Unreliable Datagram Protocol)
 - Uses IP to provide unreliable datagram delivery from process-to-process.
- TCP (Transmission Control Protocol)
 - Uses IP to provide reliable byte streams from process-to-process over connections.

Accessed via a mix of Unix file I/O and functions from the **sockets interface**.

- 14 -

15-213, F07

Hardware and Software Org of an Internet Application



- 15 -

15-213, F07

Naming and Communicating on the Global IP Internet

Original idea

- Every node on Internet would have unique IP address
 - Everyone would be able to talk directly to everyone
- No secrecy or authentication
 - Messages visible to routers and hosts on same LAN
 - Possible to forge source field in packet header

Doesn't always work that way

- We'll talk about some evolution later today

- 16 -

15-213, F07

A Programmer's View of the Internet

1. Hosts are mapped to a set of 32-bit **IP addresses**.

- 128.2.208.178

2. The set of IP addresses is mapped to a set of identifiers called Internet **domain names**.

- 128.2.208.178 is mapped to `www.cs.ohio.edu`

3. A process on one Internet host can communicate with a process on another Internet host over a **connection**.

- 17 -

15-213, F07

1. IP Addresses

32-bit IP addresses are stored in an **IP address struct**

- IP addresses are always stored in memory in network byte order (big-endian byte order)
- True in general for any words transferred in a packet header from one machine to another
 - E.g., the port numbers used in identifying an Internet connection

```
/* Internet address structure */
struct in_addr {
    unsigned int s_addr; /* network byte order (big-endian) */
};
```

Handy network byte-order conversion functions:

- `htonl`: convert long int from host to network byte order
- `htons`: convert short int from host to network byte order
- `ntohl`: convert long int from network to host byte order
- `ntohs`: convert short int from network to host byte order

- 18 -

15-213, F07