15-213

"The course that gives CMU its Zip!"

Internetworking

May 1, 2001

Topics

- Protocol layering and encapsulation
- Internetworking with hubs, bridges, and routers
- The Internet Protocol (IP)
- The global Internet

Typical computer system

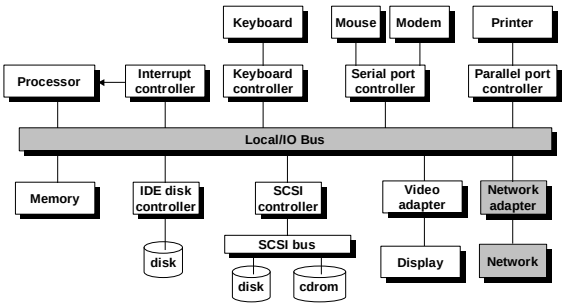

Generic network

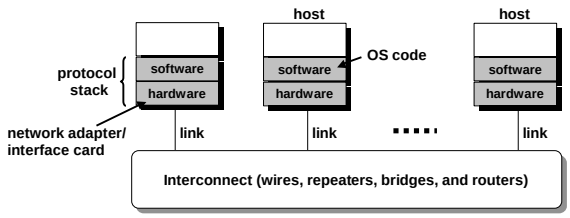

Protocols

A *protocol* defines the format of packets and the rules for communicating them across the network.

Different protocols provide different levels of service:

- simple error correction (ethernet)
- uniform name space, unreliable best-effort datagrams (host-host) (IP)
- reliable byte streams (TCP)
- unreliable best-effort datagrams (process-process) (UDP)
- multimedia data retrieval (HTTP)

Crucial idea: protocols leverage off of the capabilities of other protocols.

Protocol layering

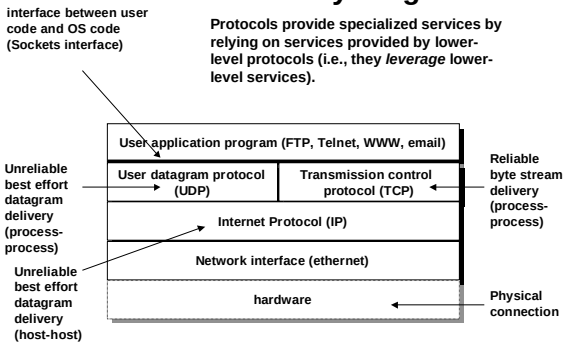

Encapsulation

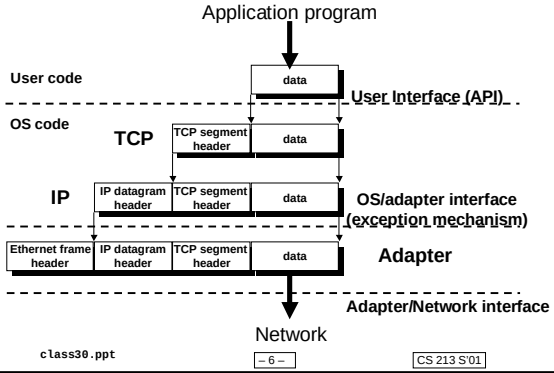

Basic network types

System area network (SAN)
- same room (meters)
- 300 MB/s Cray T3E

Local area network (LAN)
- same bldg or campus (kilometers)
- 10 Mb/sEthernet
- 100 Mb/s Fast Ethernet
- 100 Mb/s FDDI
- 150 Mb/s OC-3 ATM
- 622 Mb/s OC-12 ATM

Metropolitan area network (MAN)
- same city (10's of kilometers)
- 800 Mb/s Gigabit Nectar

Wide area network (WAN)
- nationwide or worldwide (1000's of kilometers)
- telephone system
- 1.544 Mb/s T1 carrier
- 44.736 Mb/s T3 carrier
- Global Internet

The internetworking idea (Kahn, 1972)

Build a single network (an interconnected set of networks, or *internetwork*, or *internet*) out of a large collection of separate networks.
- Each network must stand on its own, with no internal changes allowed to connect to the internet.
- Communications should be on a best-effort basis.
- "black boxes" (later called routers) should be used to connect the networks.
- No global control at the operations level.

Internetworking challenges

Challenges:

- heterogeneity
  - lots of different kinds of networks (Ethernet, FDDI, ATM, wireless, point-to-point)
  - how to unify this hodgepodge?
- scale
  - how to provide uniques names for potentially billions of nodes? (naming)
  - how to find all these nodes? (forwarding and routing)

Note: *internet* refers to a general idea, *Internet* refers to a particular implementation of that idea (The global IP Internet).

Internetworking with repeaters

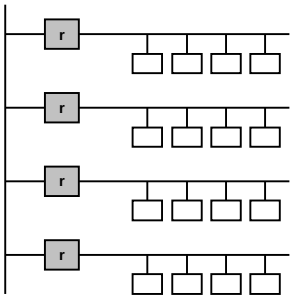

Repeaters (also called hubs) (r in the figure) directly transfer bits from their inputs to their outputs

Internetworking with repeaters

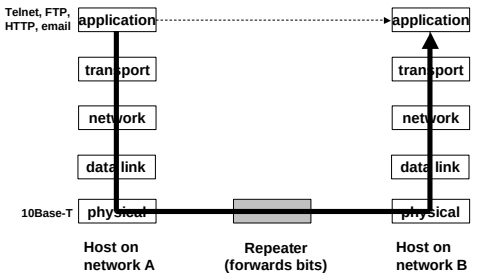

Internetworking with repeaters: Pros and cons

Pros

- Transparency
  - LANS can be connected without any awareness from the hosts.
- Useful for serving multiple machines in an office from one ethernet outlet.

Cons

- Not scalable
  - ethernet standard allows only 4 repeaters.
  - more than 4 would introduce delays that would break contention detection.
- No heterogeneity
  - Networks connected with repeaters must have identical electrical properties.

Internetworking with bridges

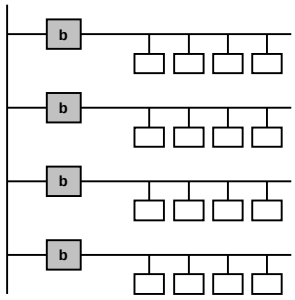

Bridges (b In the figure) maintain a cache of hosts on their input segments.

Selectively transfer ethernet frames from their inputs to their outputs.

Internetworking with bridges

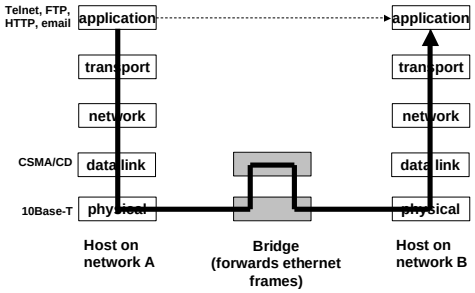

Internetworking with bridges: Pros and cons

Pros

- Transparency
  - LANS can be connected without any awareness from the hosts
  - popular solution for campus-size networks

Cons

- Transparency can be misleading
  - looks like a single Ethernet segment, but really isn't
  - packets can be dropped, latencies vary
- Homogeneity
  - can only support networks with identical frame headers (e.g., Ethernet/FDDI)
  - however, can connect different speed Ethernets
- Scalability
  - tens of networks only
    - » bridges forward all broadcast frames
    - » increased latency

Internetworking with routers

Def: An *internetwork* (internet for short) is an arbitrary collection of *physical networks* interconnected by routers to provide some sort of host-to-host packet delivery service.

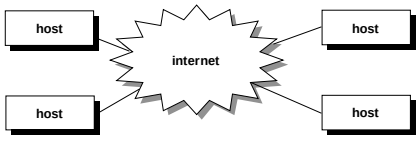

Building an internet

We start with two separate, unconnected computer networks (subnets), which are at different locations, and possibly built by different vendors.

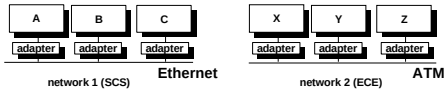

Question: How to present the illusion of one network?

Building an internet (cont)

Next we physically connect one of the computers, called a router (in this case computer C), to each of the networks.

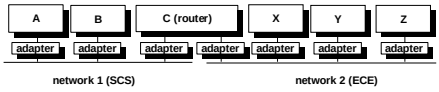

Building an internet (cont)

Finally, we run a software implementation of the Internet Protocol (IP) on each host and router. IP provides a global name space for the hosts, routing messages between network1 and network 2 if necessary.

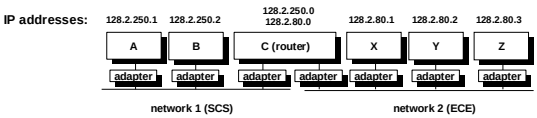

Building an internet (cont)

At this point we have an internet consisting of 6 computers built from 2 original networks. Each computer on our internet can communicate with any other computer. IP provides the illusion that there is just one network.

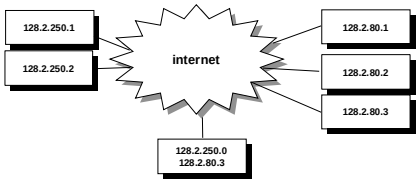

Internetworking with routers

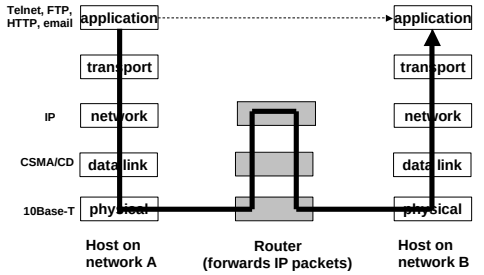

IP: Internetworking with routers

IP is the most successful protocol ever developed

Keys to success:

- **simple enough to implement on top of any physical network**
  - **e.g., two tin cans and a string.**
- **rich enough to serve as the base for implementations of more complicated protocols and applications.**
  - **The IP designers never dreamed of something like the Web.**
- **"rough consensus and working code"**
  - **resulted in solid implementable specs.**

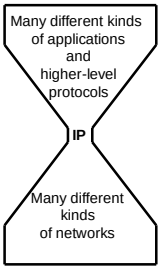

The "Hourglass Model", Dave Clark, MIT

Internet protocol stack

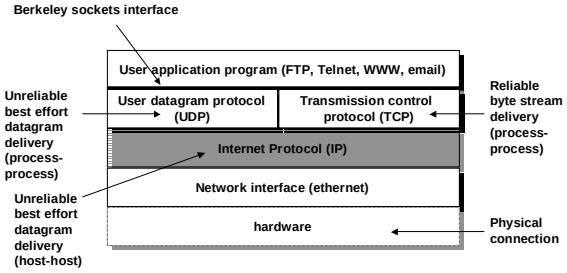

IP service model

IP service model:

- **Delivery model: IP provides best-effort delivery of datagram (connectionless) packets between two hosts.**
  - **IP tries but doesn't guarantee that packets will arrive (best effort)**
  - **packets can be lost or duplicated (unreliable)**
  - **ordering of datagrams not guaranteed (connectionless)**
- **Naming scheme: IP provides a unique address (name) for each host in the Internet.**

Why would such a limited delivery model be useful?

- **simple, so it runs on any kind of network**
- **provides a basis for building more sophisticated and user-friendly protocols like TCP and UDP**

IP datagram delivery: Example internet

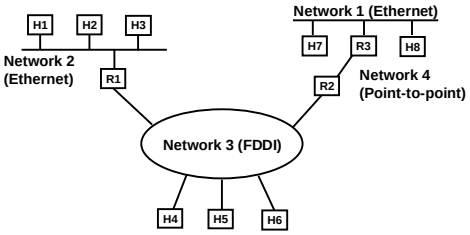

IP layering

Protocol layers used to connect host H1 to host H8 in example internet.

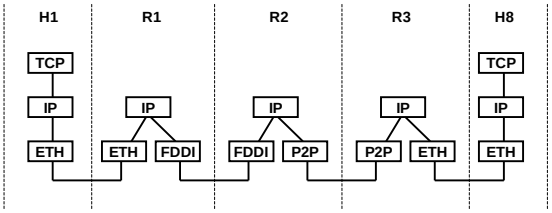

Basic Internet components

An Internet *backbone* is a collection of routers (nationwide or worldwide) connected by high-speed point-to-point networks.

A *Network Access Point (NAP)* is a router that connects multiple backbones (sometimes referred to as *peers*).

Regional networks are smaller backbones that cover smaller geographical areas (e.g., cities or states)

A *point of presence (POP)* is a machine that is connected to the Internet.

Internet Service Providers (ISPs) provide dial-up or direct access to POPs.

The Internet circa 1993

In 1993, the Internet consisted of one backbone (NSFNET) that connected 13 sites via 45 Mbs T3 links.

- Merit (Univ of Mich), NCSA (Illinois), Cornell Theory Center, Pittsburgh Supercomputing Center, San Diego Supercomputing Center, John von Neumann Center (Princeton), BARRNet (Palo Alto), MidNet (Lincoln, NE), WestNet (Salt Lake City), NorthwestNet (Seattle), SESQUINET (Rice), SURANET (Georgia Tech).

Connecting to the Internet involved connecting one of your routers to a router at a backbone site, or to a regional network that was already connected to the backbone.

The Internet backbone (circa 1993)

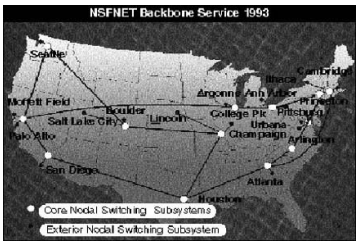

Current NAP-based Internet architecture

- In the early 90's commercial outfits were building their own high-speed backbones, connecting to NSFNET, and selling access to their POPs to companies, ISPs, and individuals.
- In 1995, NSF decommissioned NSFNET, and fostered creation of a collection of NAPs to connect the commercial backbones.
- Currently in the US there are about 50 commercial backbones connected by ~12 NAPs (peering points).
- Similar architecture worldwide connects national networks to the Internet.

Internet connection hierarchy

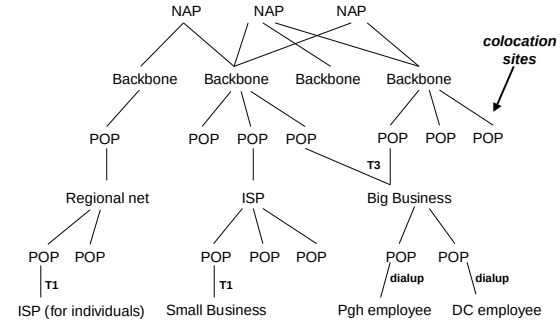

Network access points (NAPs)

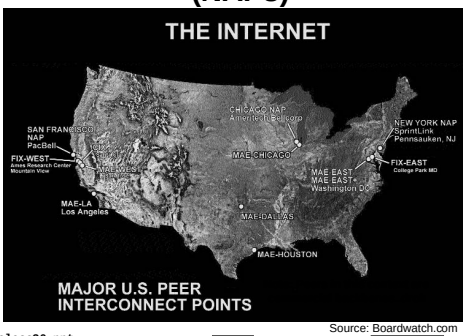

Source: Boardwatch.com

MCI/WorldCom/UUNET Global Backbone

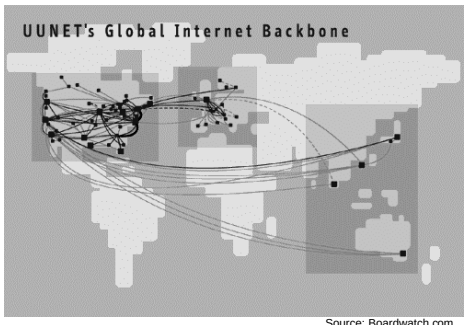

Source: Boardwatch.com