Introduction

What is a computer network?

Components of a computer network:

- hosts (nodes, computers)
- links
- protocols
- applications (network services)

Hosts & links form the *hardware* side.

Protocols & applications form the *software* side.

Protocols can be viewed as the "glue" that binds everything else together.

Simplest instance of networking problem:

> Given two hosts A, B interconnected by some network N, facilitate communication between A & B.

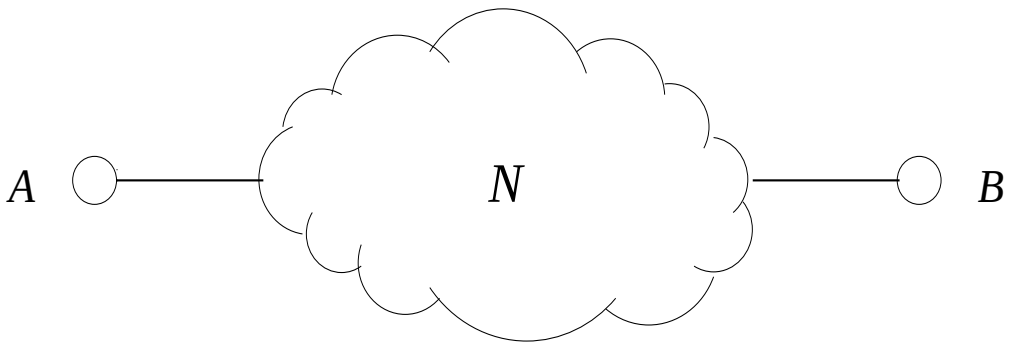

Network N can be

- point-to-point links
- multi-access links
- internetworks

Point-to-point links

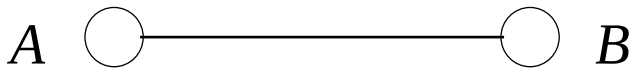

$\longrightarrow$ no addressing necessary

Multi-access links

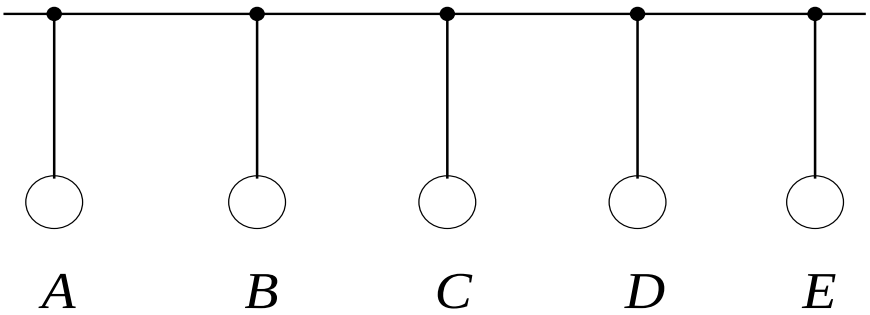

$\longrightarrow$ addressing necessary

$\longrightarrow$ access control ("bus arbitration")

Internetwork

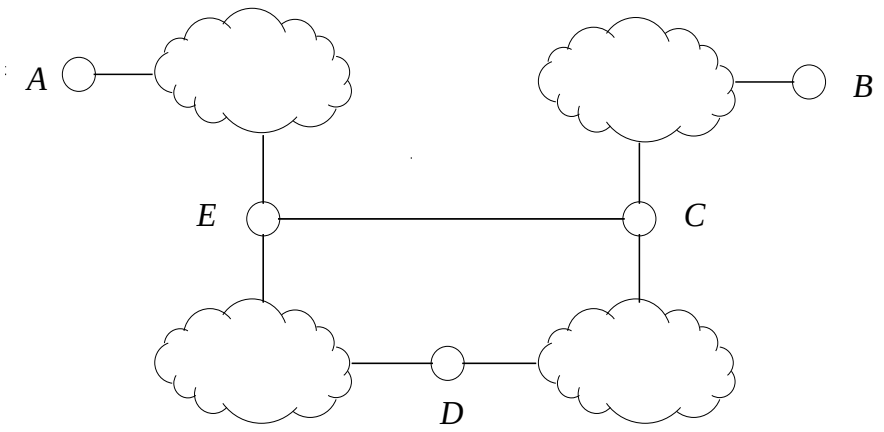

$\longrightarrow$ addressing necessary

$\longrightarrow$ path selection (routing)

$\longrightarrow$ protocol translation (internetworking)

Myriad of different local area network (LAN) technologies co-existing in a wide area network (WAN).

e.g.: Ethernet, token bus, FDDI (Fiber Distributed Data Interface), packet radio, etc.

Notice that communicating entities are *processes* residing on A and B.

Thus an *address* must also identify which process a message is destined for.

Fault-tolerance issues.

$\longrightarrow$ the larger the network, the more things can go wrong.

e.g.: node/link failures, message corruption, lost messages, outdated messages.

$\longrightarrow$ issue of reliable communication.

Performance issues.

$\longrightarrow$ potential for bottleneck development.

Network performance

Two yardsticks (performance measures):

$\longrightarrow$ throughput, delay (latency).

Bandwidth vs. throughput:

Bandwidth—maximum data transmission rate achievable at the hardware level; determined by signalling rate of network interface card (NIC) and physical link.

Throughput (*effective*)—maximum data transmission rate achievable at the software level; overhead of network protocols is accounted for.

$\longrightarrow$ "true" measure of communication speed.

Mbps (Mb/s), kbps (kb/s)

$\longrightarrow$ 10^6 bits per second, 10^3 bits per second; indicates data transmission rate; influenced by clock rate (MHz) of signalling hardware.

e.g.: 10 Mbps (10BaseT), 100 Mbps (FDDI), 1.544 Mbps (T1), 44.736 Mbps (T3), 155.52 Mbps (OC-3), 622.08 Mbps (OC-12), etc.

MB, kB

$\longrightarrow$ 2^{20} bytes, 2^{10} bytes; indicates size of data being shipped; influenced by the memory structure of computer.

e.g.: 512 B, 1 kB (TCP segment size), 64 kB (maximum IP packet size), 53 B (ATM cell), 810 B (SONET frame), etc.

Packet, frame, cell, datagram, message, etc.

$\longrightarrow$ *packet* most generic term (unit of data).

Message delay (latency):

Given a message M of size s, its *delay* or *latency*, $d(M)$, is given by

$$d(M) = D_P + D_T(M) + D_Q(M, \mathbf{S})$$

where D_P is the propagation time, $D_T(M)$ is the transmission time, $D_Q(M, \mathbf{S})$ is the queueing delay, and $\mathbf{S}$ is the network state.

Furthermore,

$$D_P \propto L/c$$

where L is the total distance from source to destination and c is the speed of light.

$\longrightarrow$ D_P is independent of M.

$$D_T(M) \propto s/BW$$

where BW is the link bandwidth.

Queueing delay, $D_Q(M, \mathbf{S})$, generally, does not have a closed-form expression, being a complicated function of $\mathbf{S}$.

$\longrightarrow$ in fact, it is the most important factor.

$*$ *Implicit in the previous formulae: infinite (suff. large) buffer size. Why?*

Lesson from queueing theory:

> Other things being "equal," if traffic flow is *stochastic* (with certain properties), then there exists a trade-off relationship between throughput and delay.

Definition of *utilization* ρ:

$$\rho = \text{throughput}/BW.$$

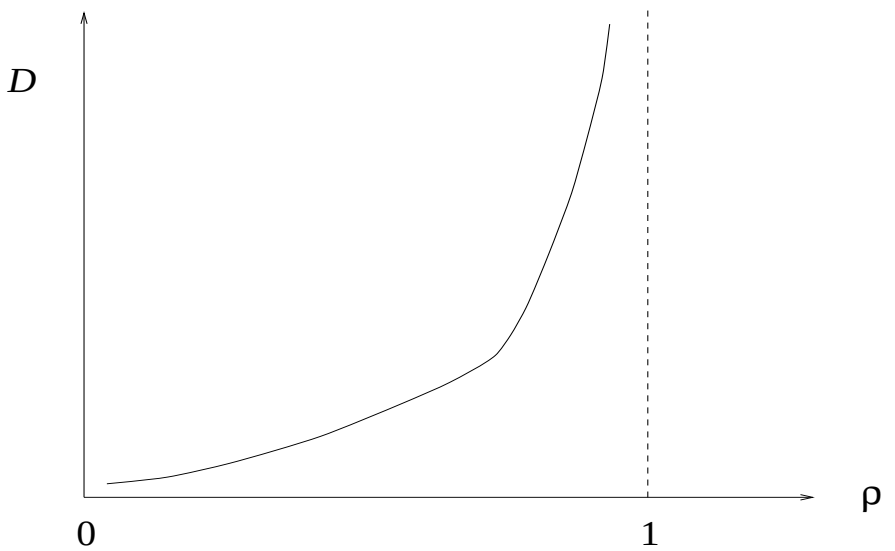

Why is this relationship important?

> *Multi-media data* (voice, image, real-time interactive data) needs to meet certain quality-of-service (QoS) requirements to be useful to the user.
>
> $\longrightarrow$ small delay & high utilization is desirable.
>
> Trade-off relationship points toward *fundamental* limitations in achieving both objectives at the same time.

Other performance measures:

- packet (cell) loss rate
- delay variation—"jitter"
- throughput variation—"burstiness"
- etc.

What is traveling on the wires?

Mixed data:

$\longrightarrow$ bulk data, audio, video, real-time interactive data, etc.

Tilting more and more every day toward *multi-media* data; i.e., traffic with real-time constraints.

Data networks capable of carrying all types of data at the same time is a *new* phenomenon.

Even today, *audio traffic* (telephone) is carried on an entirely separate communication network vis-a-vis *data traffic*, operating under different internetworking principles from the latter.

$\longrightarrow$ Time-division multiplexing (TDM) for telephony vs. packet switching for data networks.

Time-division multiplexing:

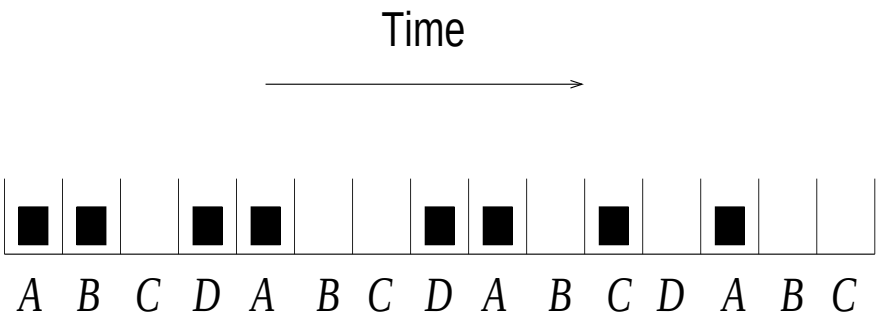

Packet/circuit switching:

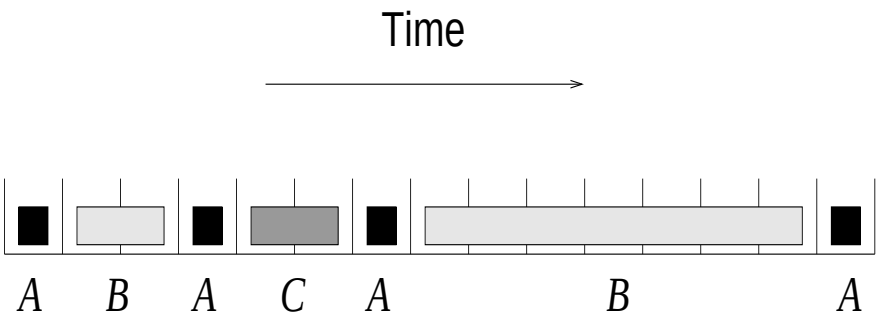

$\longrightarrow$ "Bully phenomenon." What to do?

Asynchronous transfer mode (ATM) :

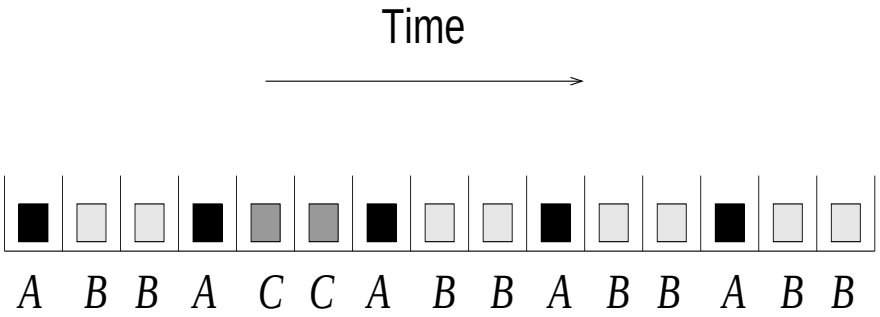

$\longrightarrow$ 53 byte packet or *cell*.

Synergy of all forms of data, audio, video, bulk, etc. One unified network with "integrated" services.

Still many problems: how to resolve conflicts? Which cell gets priority in case of contention; what is the justification? Is "fairness" preserved? Can quality-of-service (QoS) be guaranteed, if need be; and if so, at what cost?

$\longrightarrow$ *performance* question.

How to make sense of all this?

Study of networks can be divided into three aspect:

- architecture
- algorithms
- implementation

Architecture

- hardware
  - communication or data link technology (e.g., Ethernet, SONET)
  - hardware interface standards (e.g., EIA RS 232C—serial communication between DTE and DCE)
- software
  - conceptual organization (e.g., ISO/OSI 7 layer reference model, ATM network model)
  - protocol standards (e.g., IAB RFC—TCP, UDP, IP; the *what* over *how*)

. . . provides the "shell" or "skeleton" to everything else.

... speaking of *standards*,

- ISO (International Standards Organization). ISO/OSI 7-layer reference model.
- ANSI (American National Standards Institute). U.S. Government representative at ISO.
- CCITT (International Telegraph and Telephone Consultative Committee). U.N.-chartered organization, principal international telecommunication standards organization.
- ITU (International Telecommunications Union). Successor of CCITT (used to be parent organization), U.N.-chartered.
- ATM Forum. Industry organization.
- IEEE. Professional society, LAN standards; e.g., IEEE 802.3 (Ethernet), IEEE 802.5 (token ring).
- IAB (Internet Architecture Board). Informal committee reporting to DoD, NSF.

- IETF (Internet Engineering Task Force). Subsidiary of IAB; main concern—short-term issues.
- IERF (Internet Research Task Force). IAB subsidiary; long-term issues.
- Internet Society. Elected trustees appoint IAB members.

Algorithms

- error detection and correction (e.g., checksum, CRC)
- routing (e.g., shortest paths—Dijkstra, Bellman & Ford)
- congestion control or flow control (e.g., TCP window control and RTT estimation)
- traffic shaping (e.g., leaky bucket) and admission control
- queueing discipline (e.g., FCFS, weighted fair queueing)

Most important component; the *how* aspect of computer networks.

Directly impacts network performance by controlling the underlying resources provided by the network architecture.

Implementations

Same algorithm or protocol is implemented, i.e., *coded*, in several ways (e.g., TCP from different vendors).

$\longrightarrow$ network programming

Main issue: *efficiency*. In particular, speed.

- reduce copying operation.
- header prediction.
- employ multi-threading to reduce context-switch time overhead; horizontal vs. vertical organization.
- employ multi-threading to hide communication latency.

Even though at times ugly, a *must* to squeeze the best out of performance.

Layering

Most fundamental organization of the software aspect of computer networks.

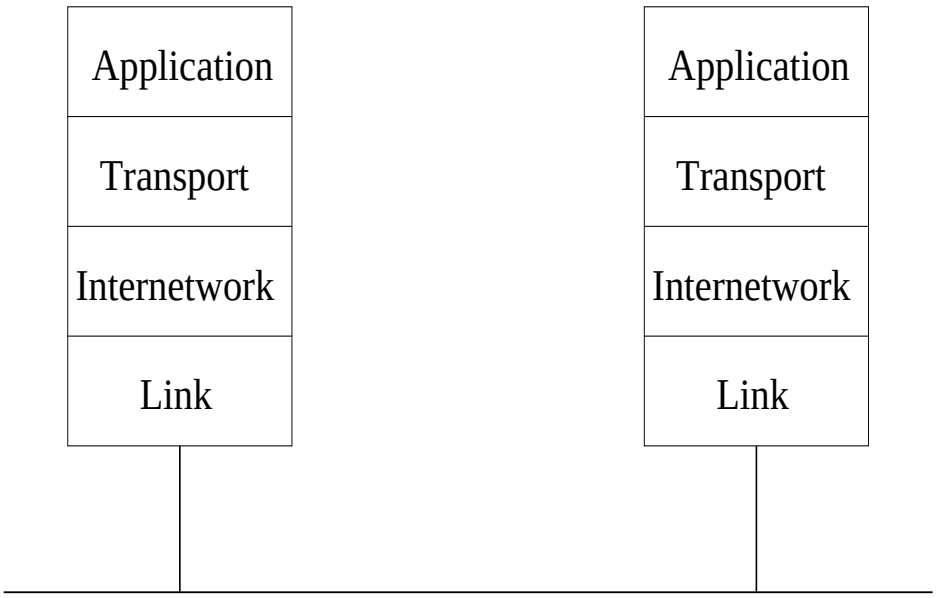

Achieves abstraction, modularization; a logical picture of computer networks incorporating the functionality at every node.

Encapsulation

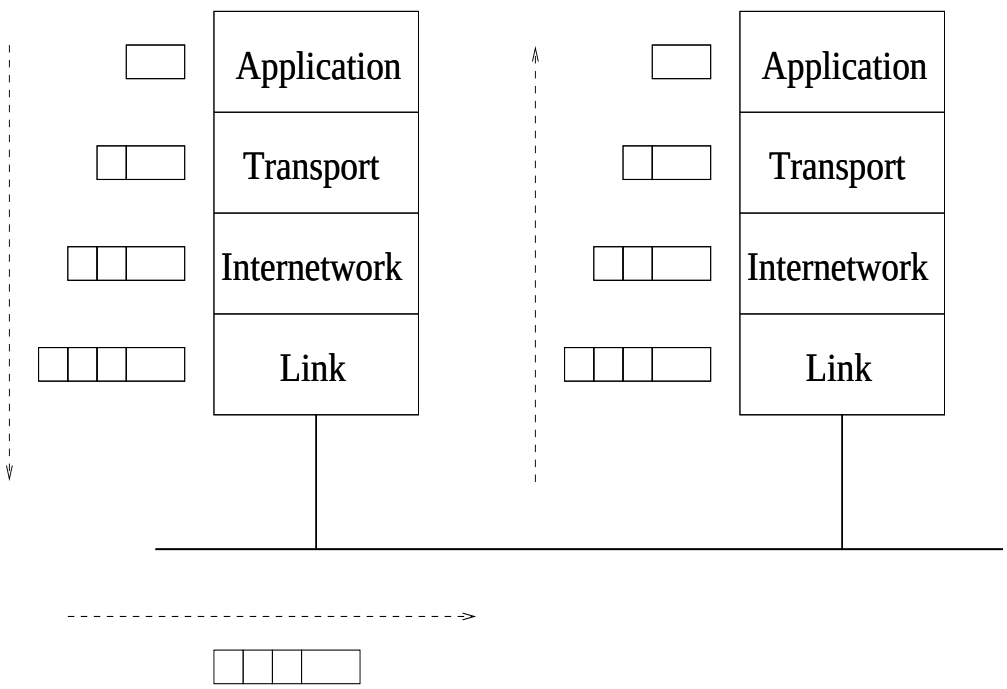

Header and trailer overhead; redundancy (extreme case: IP over ATM); efficiency issue (header prediction, horizontal or flat organization) for parallel processing.

ISO/OSI 7-layer reference model

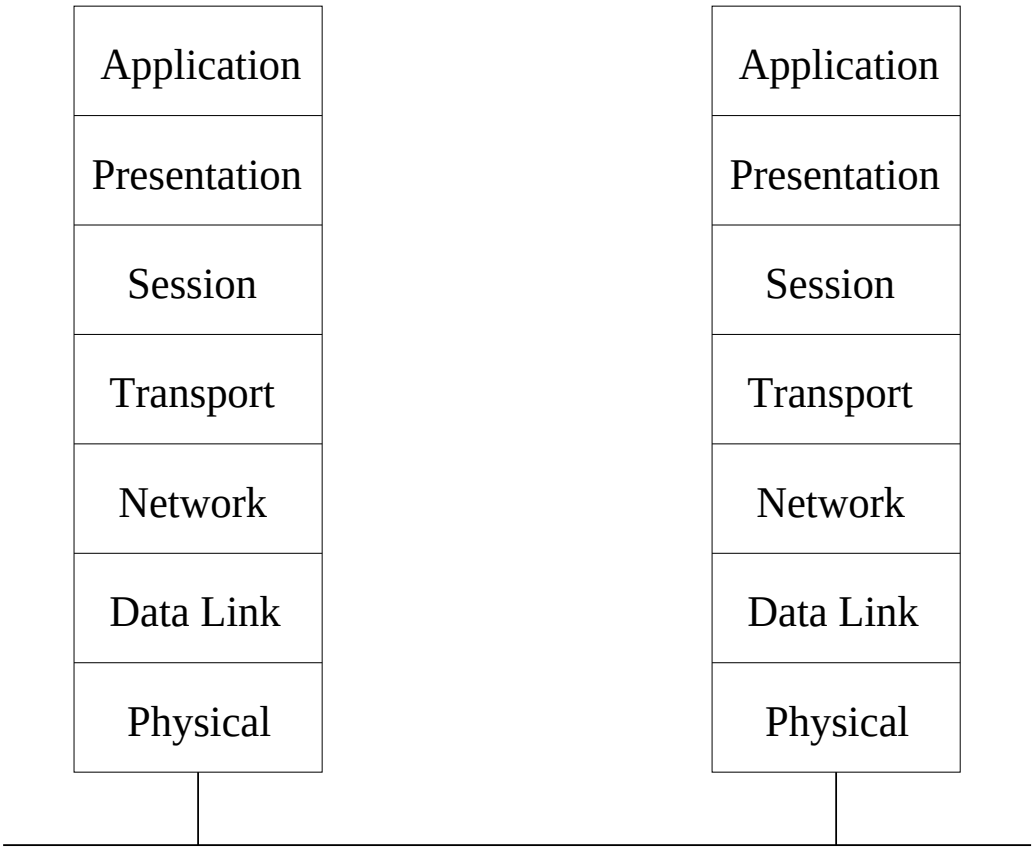

Outdated; however, still useful as logical reference point.

Protocol graph

Shows logical relationship between protocol modules in the protocol stack.

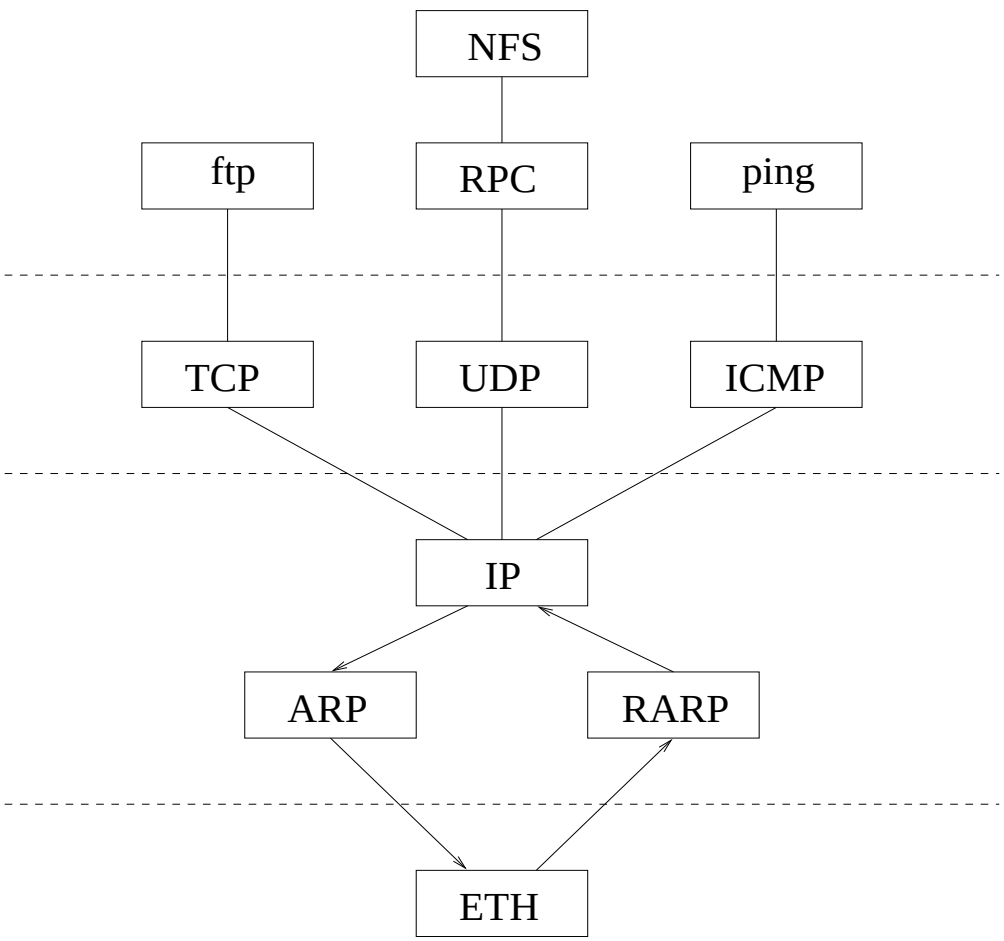

Circuit switching vs. packet switching

Circuit switching—virtual channel is established and followed during the duration of an end-to-end conversation or connection.

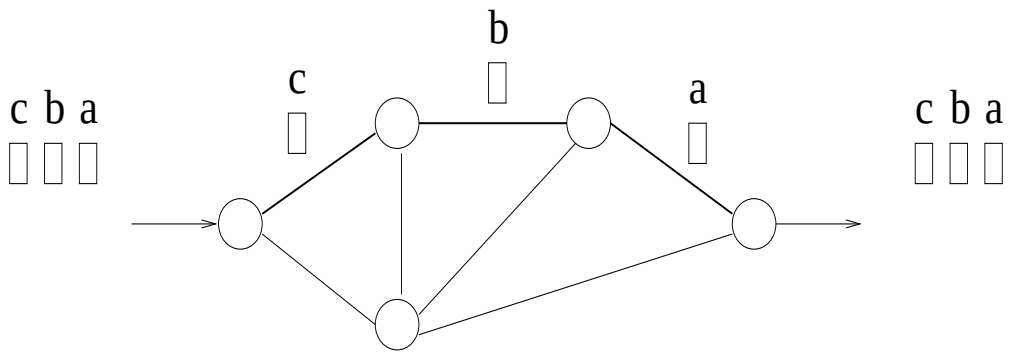

$\longrightarrow$ static route

$\longrightarrow$ in-order delivery

Telephone networks and ATM networks.

Packet switching—every packet belonging to an end-to-end conversation or connection is an independent entity and may take a different route from other packets in the same conversation.

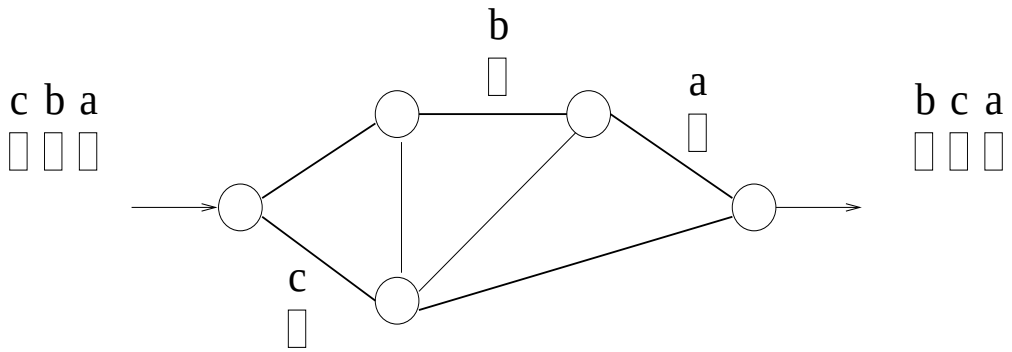

$\longrightarrow$ dynamic route

$\longrightarrow$ out-of-order delivery

Trade-off between processing overhead and route overhead; if route is fixed, the probability of the virtual circuit becoming an inefficient route increases with duration of conversation.

Reliable vs. unreliable communication

Packets may get

- corrupted due to errors (e.g., noise)
- dropped due to buffer overflow
- dropped due to aging or outdatedness—TTL (time-to-live field in IP)
- lost due to link or host failures

Internet philosophy: reliable transport (TCP) over unreliable internetwork (IP). Use retransmission and acknowledgment (ack).

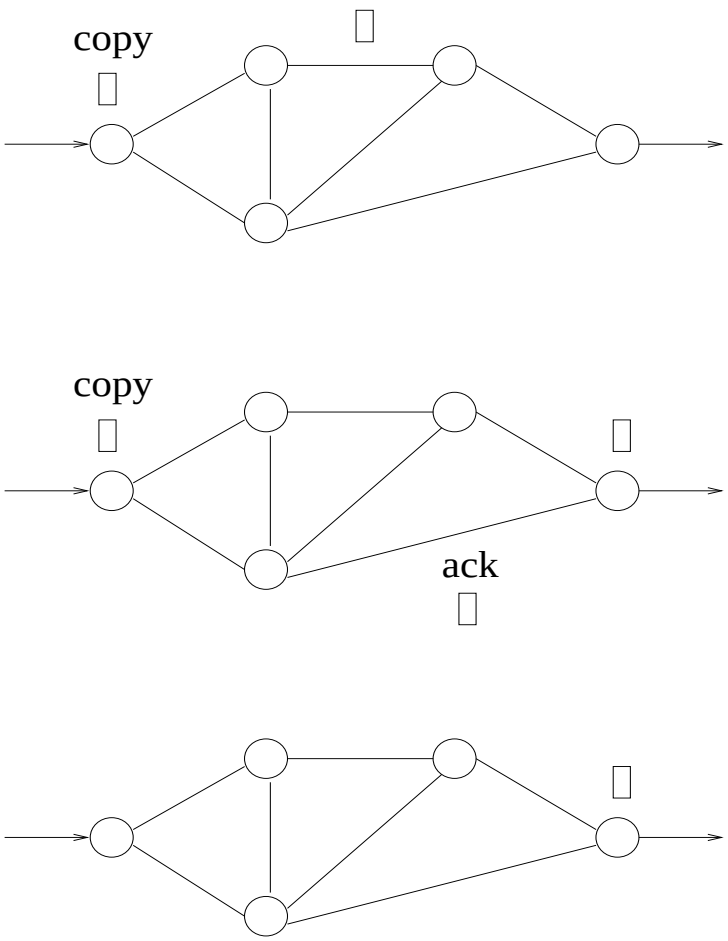

Reliability can also be implemented at the data link layer.

Trade-off between link-level overhead (fixed) and end-to-end reliability overhead (variable).

Sharing of resources

Networks can be viewed as the ultimate platform for sharing *resources*

- software resources.
  - passive information—images, sound, text, etc.
  - active information—programs, network computing (e.g., Java)
- hardware resources.
  - system resources—printers, disk, CPU, tape, etc.
  - network resources—bandwidth, buffer space, CPU

Too much sharing: *resource contention.*

Resolution and prevention of resource contention is a primary task of software layer.

Queueing and buffer overflow

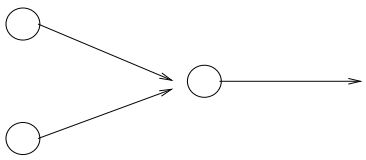

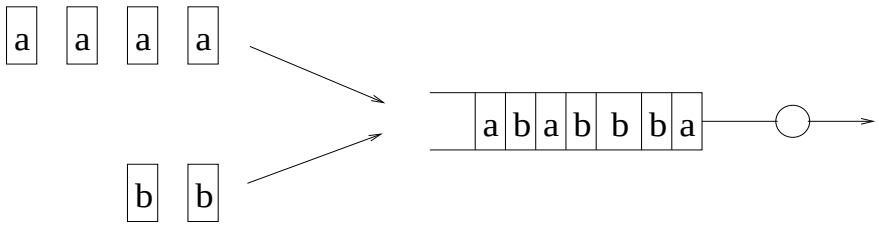

Queueing or scheduling policy

- dequeueing policy (e.g., FCFS)
- enqueueing policy (reshuffling, who-to-drop, overwrite)

Based on the notion of *priorities* (e.g., UNIX scheduling of processes).

IP has TOS (type-of-service) but not used; IPv6 (Next Generation IP) actively implements TOS.

One way to deal with quality-of-service issue

$\longrightarrow$ differential treatment of packets

Quality-of-service (QoS) is one of the main issues in ATM networks.

Resource allocation problem

Various forms of optimization problems

- network flow problem (e.g., routing and multi-commodity flow, routing and shortest path)
- queueing theory and dynamical systems (e.g., congestion control, QoS)

Issues: seek optimal *operating point* of the system while maintaining stability.

Stability is a big issue (e.g., routing, congestion control)

One definition of optimal operating point: *power*.

$$\text{power} = \frac{\text{throughput}}{\text{delay}}.$$

Trade-off measure capturing throughput normalized by delay.