Foundation

CSCI 363 Computer Networks

Department of Computer Science

Textbook: *Computer Networks, A Systems Approach*
by Larry Peterson and Bruce Davie

What is this *systems approach*?

- “Start with first principles and walk through the thought process that let to today’s networks.”
- The book follows a bottom-up approach, while we follow a top-down in this class.
- Discuss real protocols, rather than just abstractions.
- Stronger emphasis on **system software**: networking hardware is commodity off-the-shelf. It is in software that new services and applications are defined.
- Networks are a complex composition of smaller building blocks. Look at end-to-end behavior not only individual components.
- Emphasis on empirical analysis.

Connectivity

Wish List:

- Interconnect machines.
- Maintain data *confidentiality*, data *integrity*, and system *accessibility*.
- Support growth by allowing more and more computers, or nodes, to join in (*scalability*).
- Support increases in geographical coverage.

Links

Each node needs one interface (NIC) for each link.

point-to-point

multiple-access

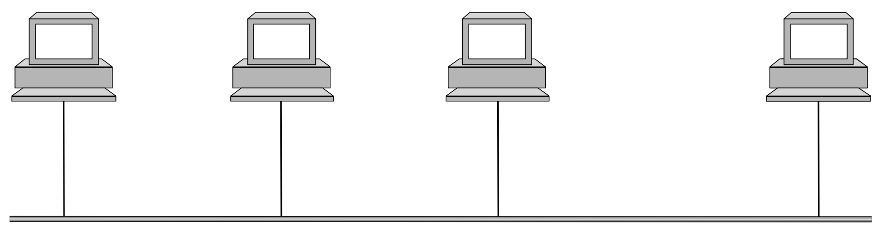

Geographical coverage and scalability are limited.

Switched Networks

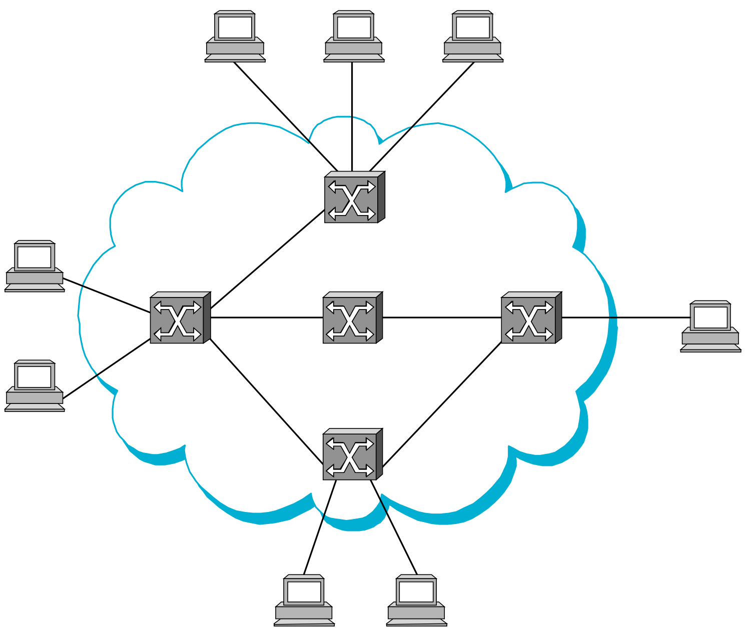

Circuit Switched

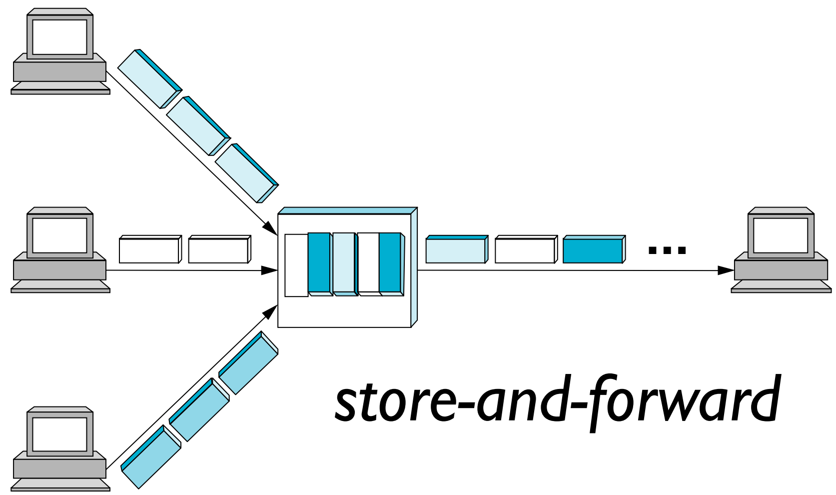

Packet Switched

Internetworking

To interconnect two or more networks, one needs a ***gateway*** or ***router***.

Host-to-host connectivity is only possible if there's a uniform ***addressing*** scheme and a ***routing*** mechanism.

Messages can be sent to a single destination (***unicast***), to multiple destinations (***multicast***), or to all possible destinations (***broadcast***).

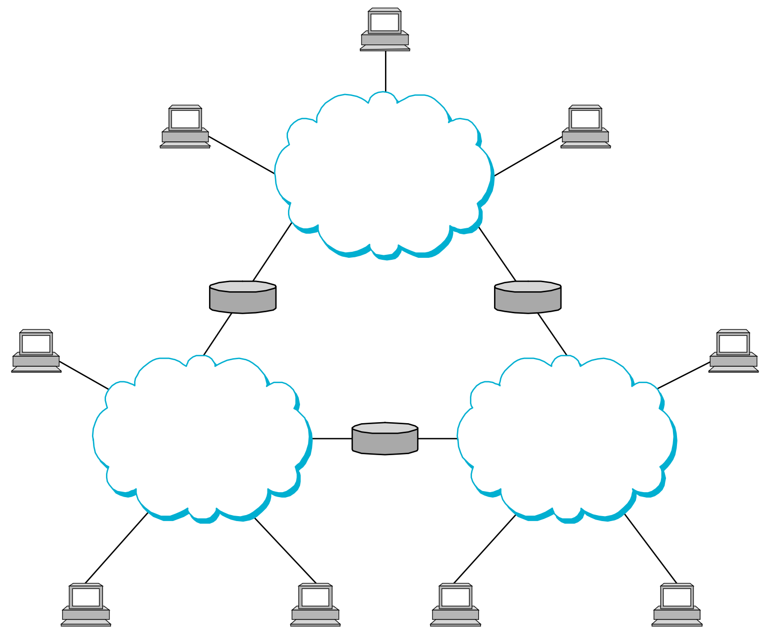

Sharing a Link

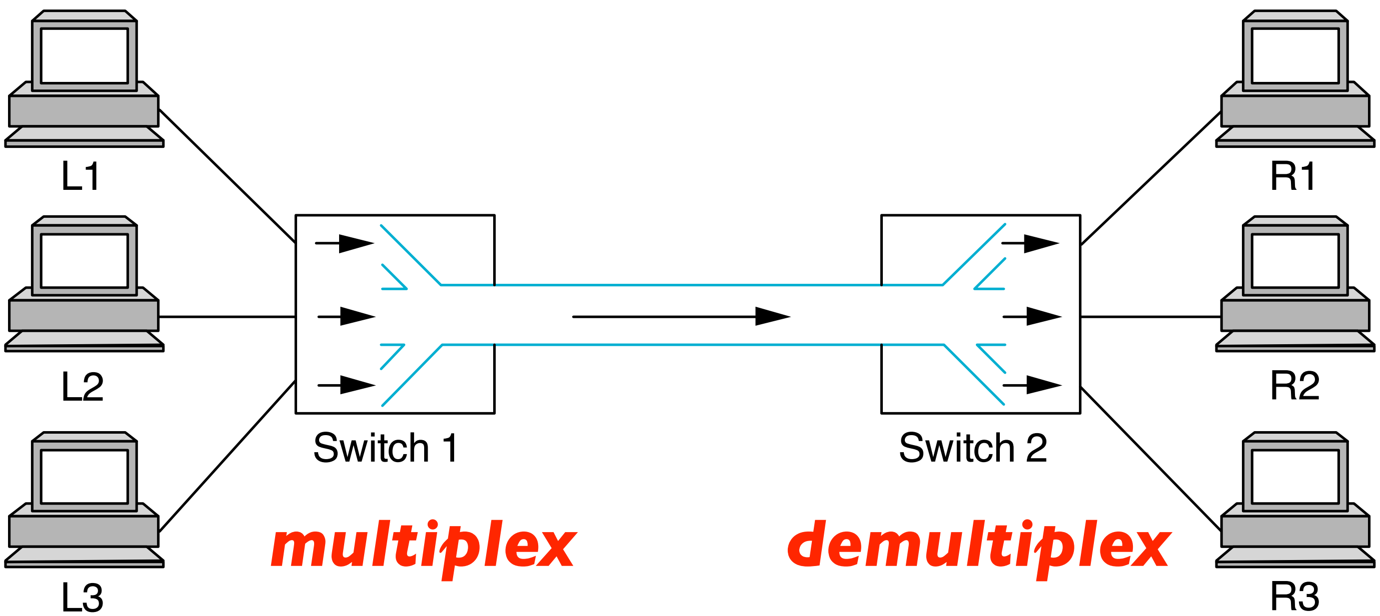

Synchronous Time Division Multiplexing (STDM)

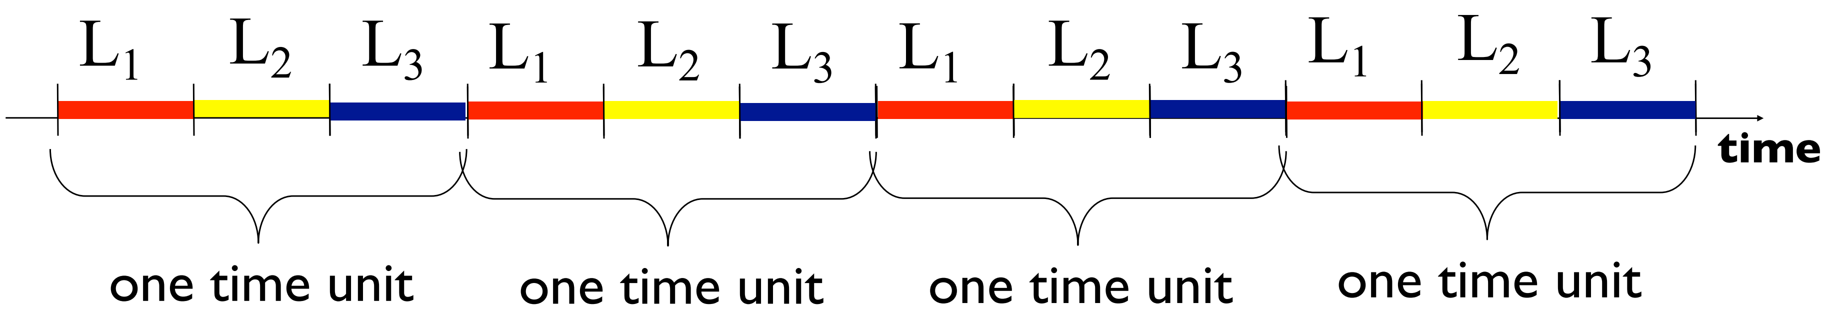

Divide time into equal-sized quanta and assign each them to flows on the physical link in round-robin fashion.

Frequency-Division Multiplexing (FDM)

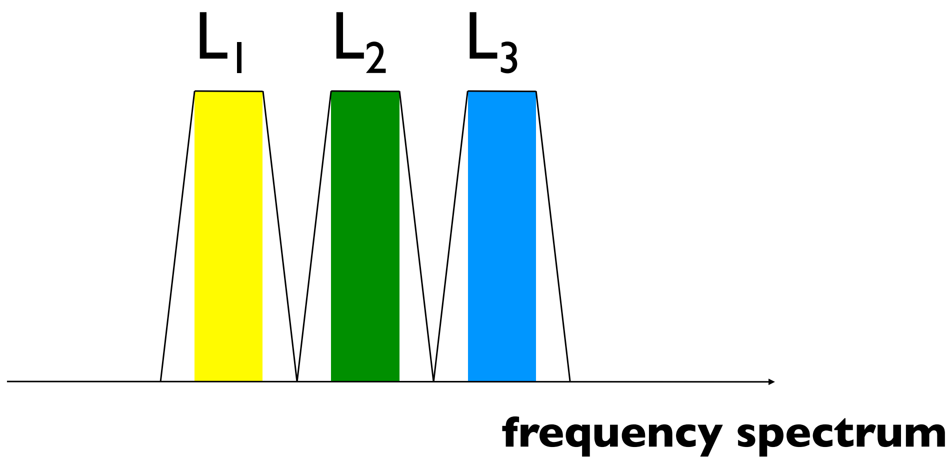

All flows are transmitted simultaneously on the link, but each one uses a different frequency.

Statistical Multiplexing

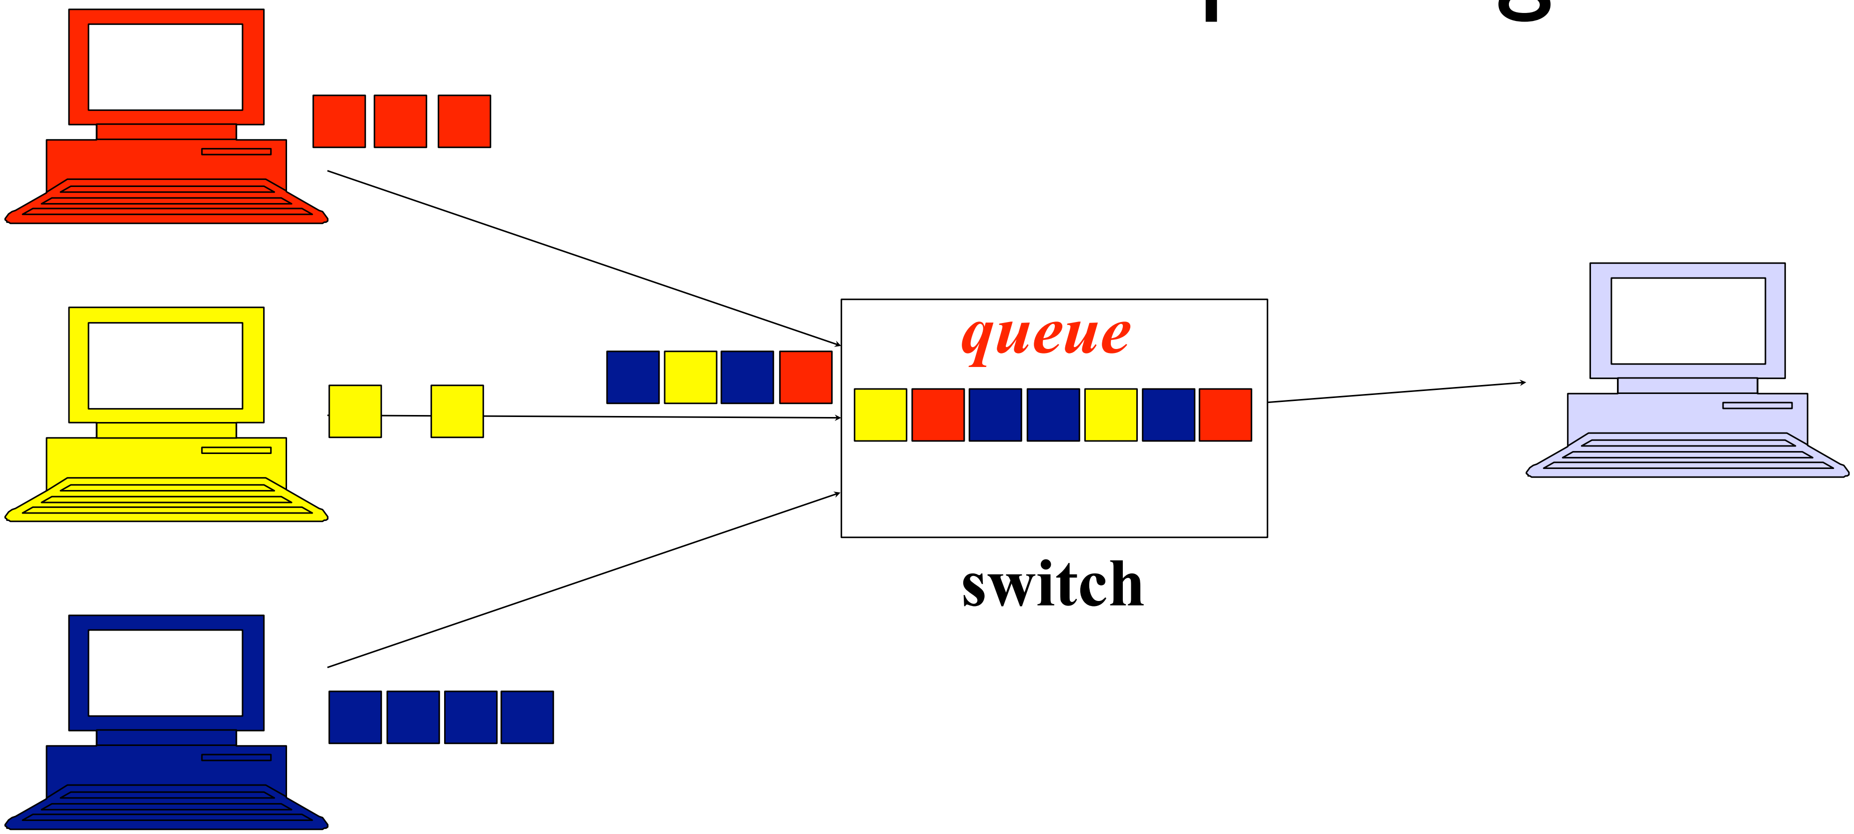

Each flow is broken into packets and sent to a switch, which can deal with the arriving packets according to a policy (FIFO, round-robin, etc).

Analytical Framework

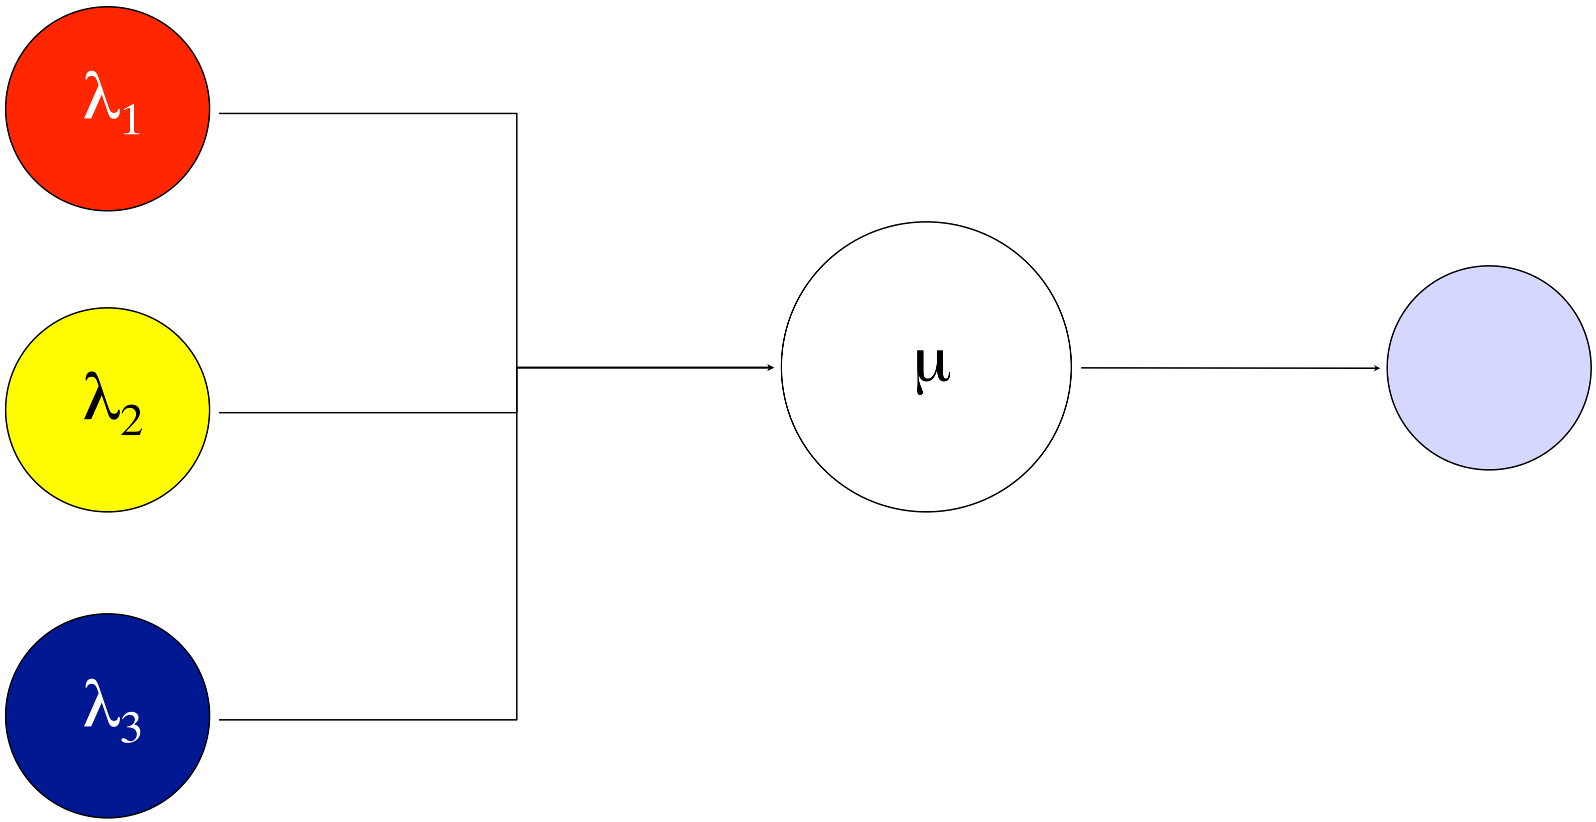

Say λ_1, λ_2, and λ_3 are *packet arrival rates* and μ the *service rate*. If we can characterize the probability distributions of packet interarrival times and of packet service times, ***queueing theory*** can help us compute metric such as throughput, wait time, etc.

Reliability

Networks must deal with:

- Physical damage to cables,
- Electromagnetic interference,
- Machine crashes and reboots,
- Memory limitations,
- Software bugs.

Classes of failure:

- ***Bit errors*** (single bit or burst),
- ***Packet loss,***
- ***Link and node failures.***

Challenge: Fill in the gap between what applications expect of the medium and what underlying technologies can actually provide.

Range of Coverage

We can classify computer networks according to their geographical coverage:

LAN: local area network.

WLAN: wireless local area network.

MAN: metropolitan area network.

WAN: wide area network (long haul network).

In interconnecting multiple networks (internetworking), we're interested in the seamless integration of all these levels. Note that different levels use very different technologies.

Network Architecture

Application programs
Process-to-process channels
Host-to-host connectivity
Hardware

How is a layered architecture helpful in the design of networks that meet the goals we stated?

The ISO/OSI Reference Model

Source: Computer Networks, Andrew Tanenbaum

ISO: International Standards Organization
OSI: Open Systems Interconnection

The protocol *stack*:

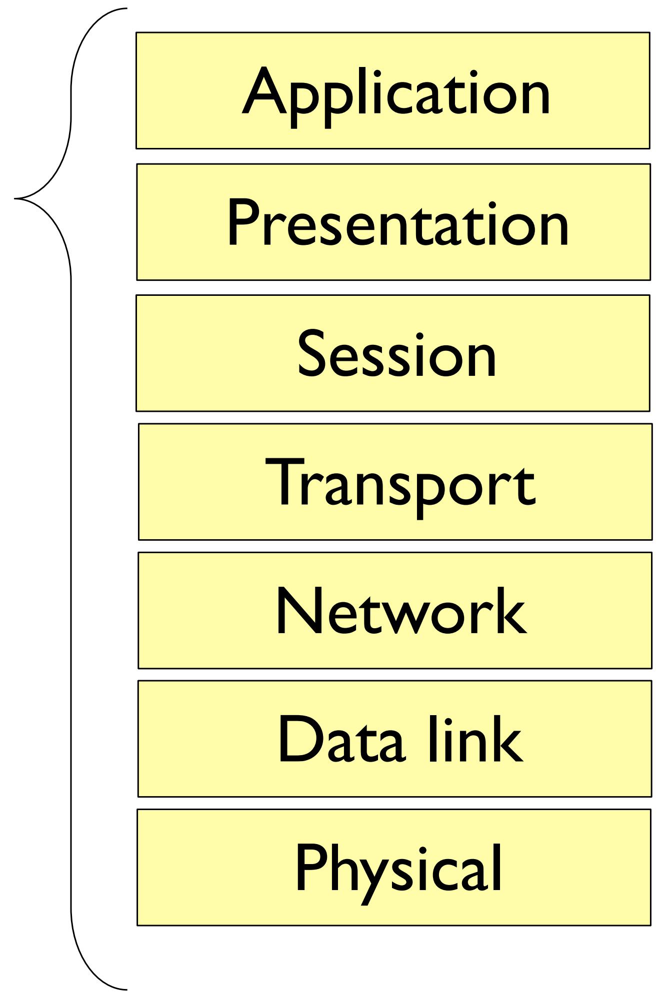

The idea behind the model: Break up the design to make implementation simpler. Each layer has a well-defined function. Layers pass to one another only the information that is relevant at each level. Communication happens only between adjacent layers.

The Layers in the ISO/OSI RF Model

Physical: Transmit raw bits over the medium.

Data Link: Implements the abstraction of an error free medium (handle losses, duplication, errors, flow control).

Network: Routing.

Transport: Break up data into chunks, send them down the protocol stack, receive chunks, put them in the right order, pass them up.

Session: Establish connections between different users and different hosts.

Presentation: Handle syntax and semantics of the info, such as encoding, encrypting.

Application: Protocols commonly needed by applications (cddb, http, ftp, telnet, etc).

Communication Between Layers within a Host

It’s important to specify the services offered to higher layers in the hierarchy. What they are + how to use them = *interface*.

SAPs (service access points)

<u>Note:</u> This is ISO terminology.

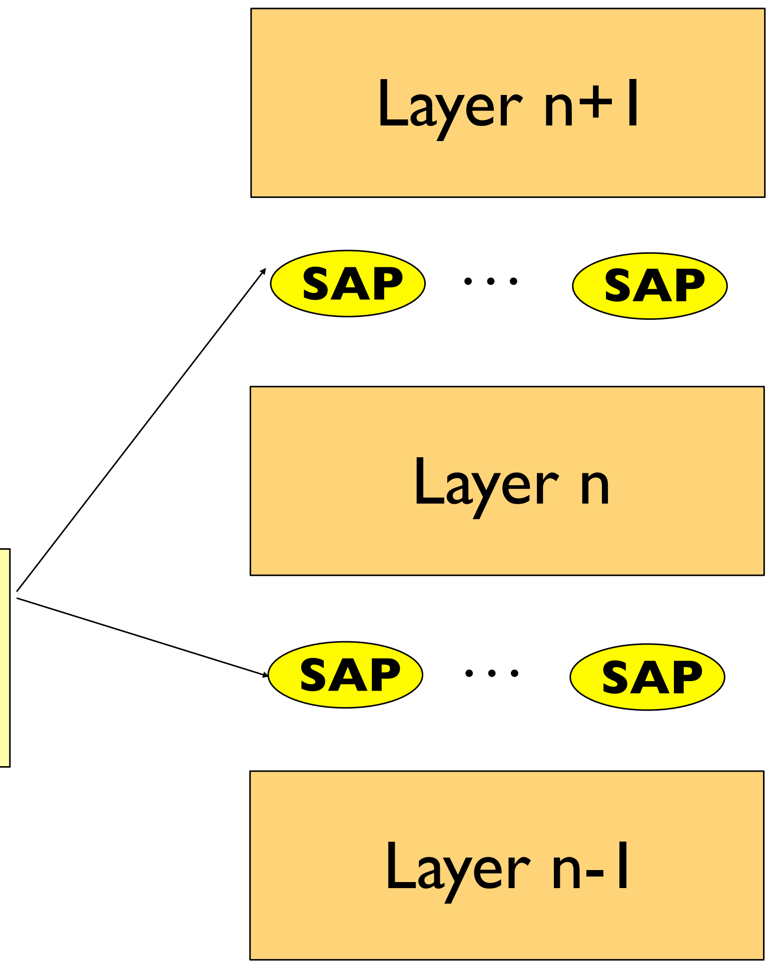

Communication Between Layers in Different Hosts

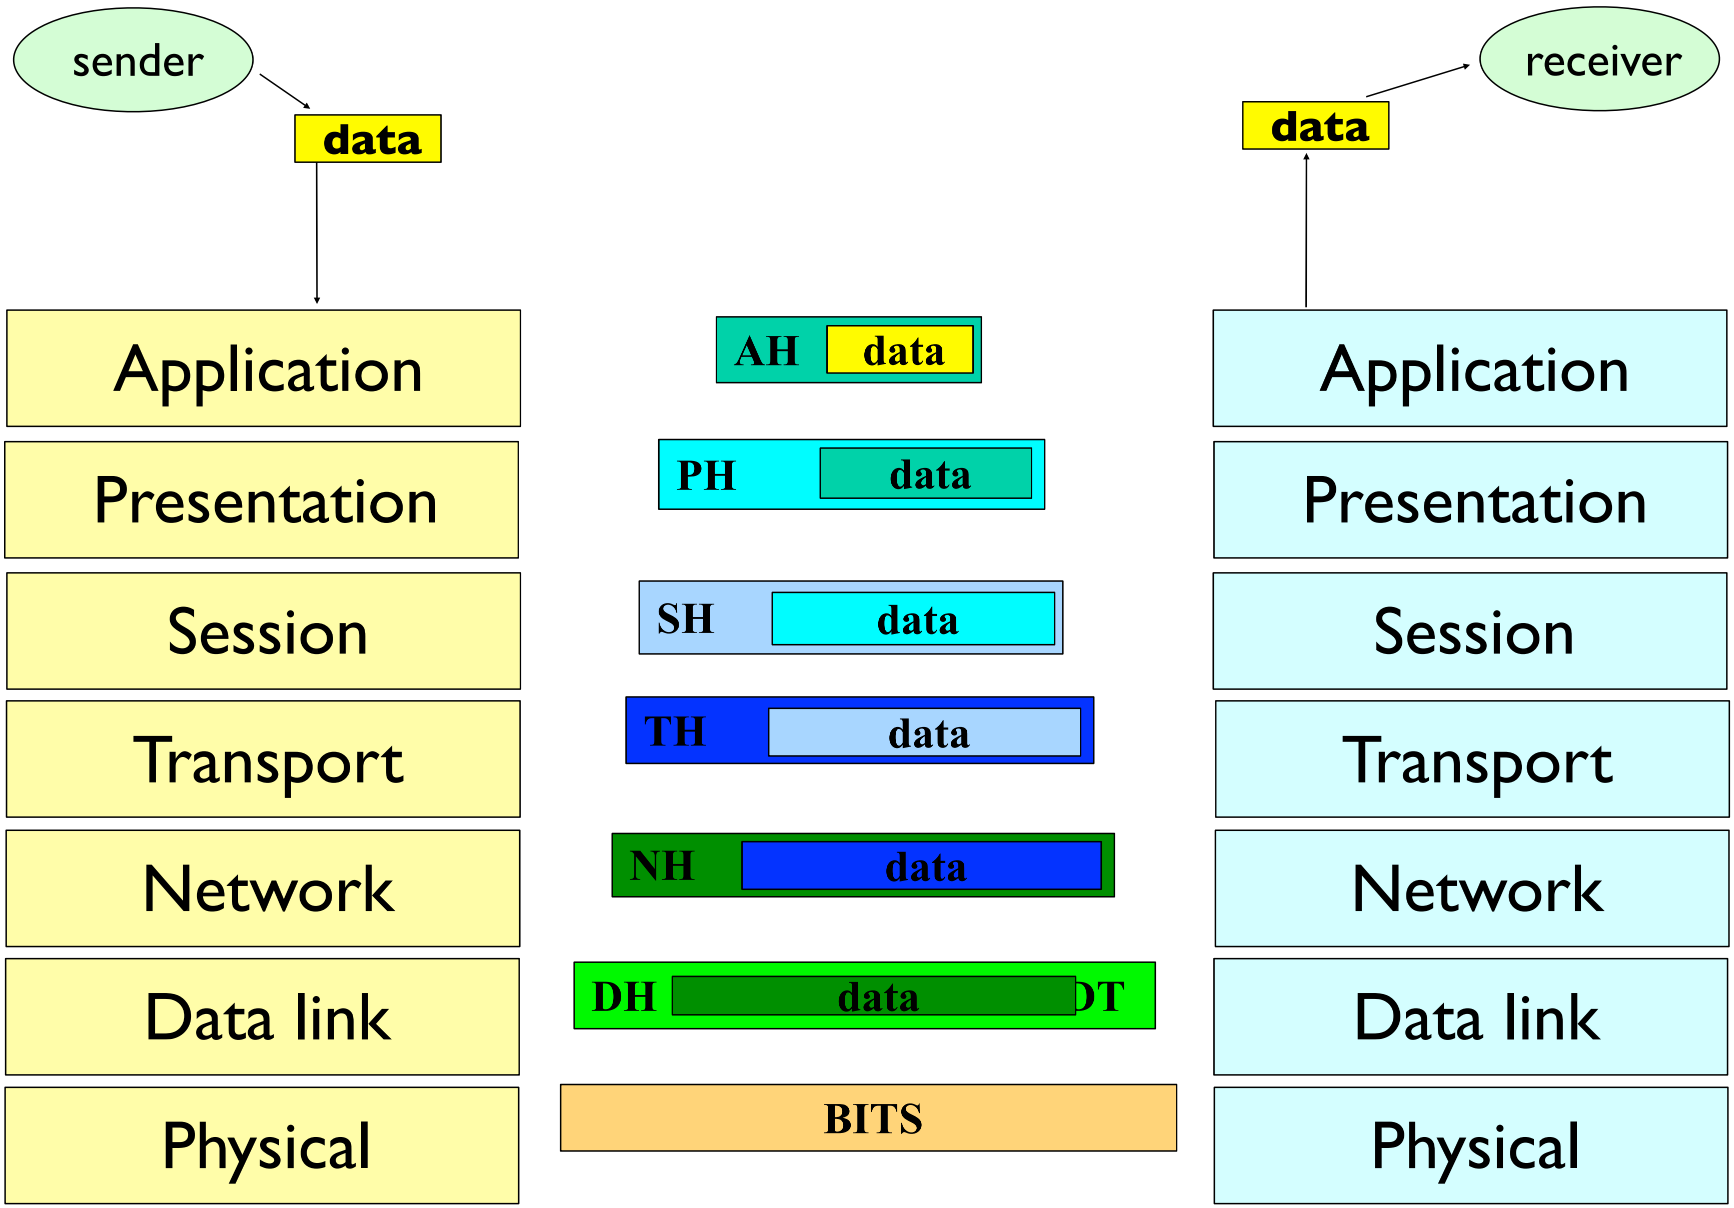

Communication Between Layers in Different Hosts

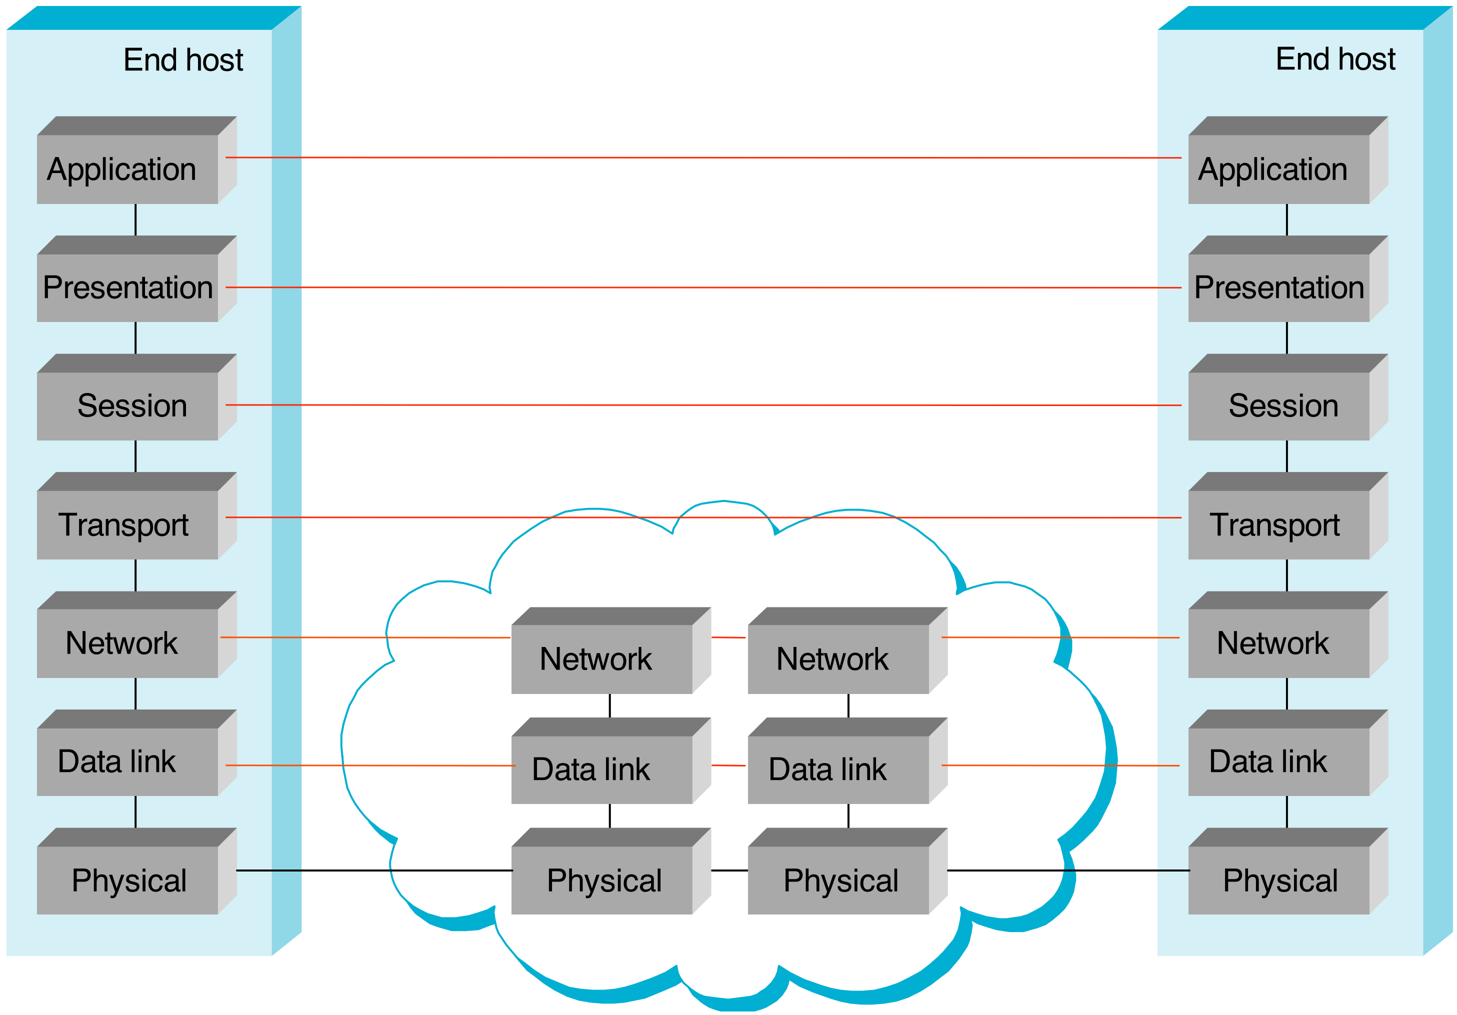

One or more nodes within the network.

The Layers in the TCP/IP Protocol Suite

Source: The TCP/IP Protocol Suite, Behrouz A. Forouzan

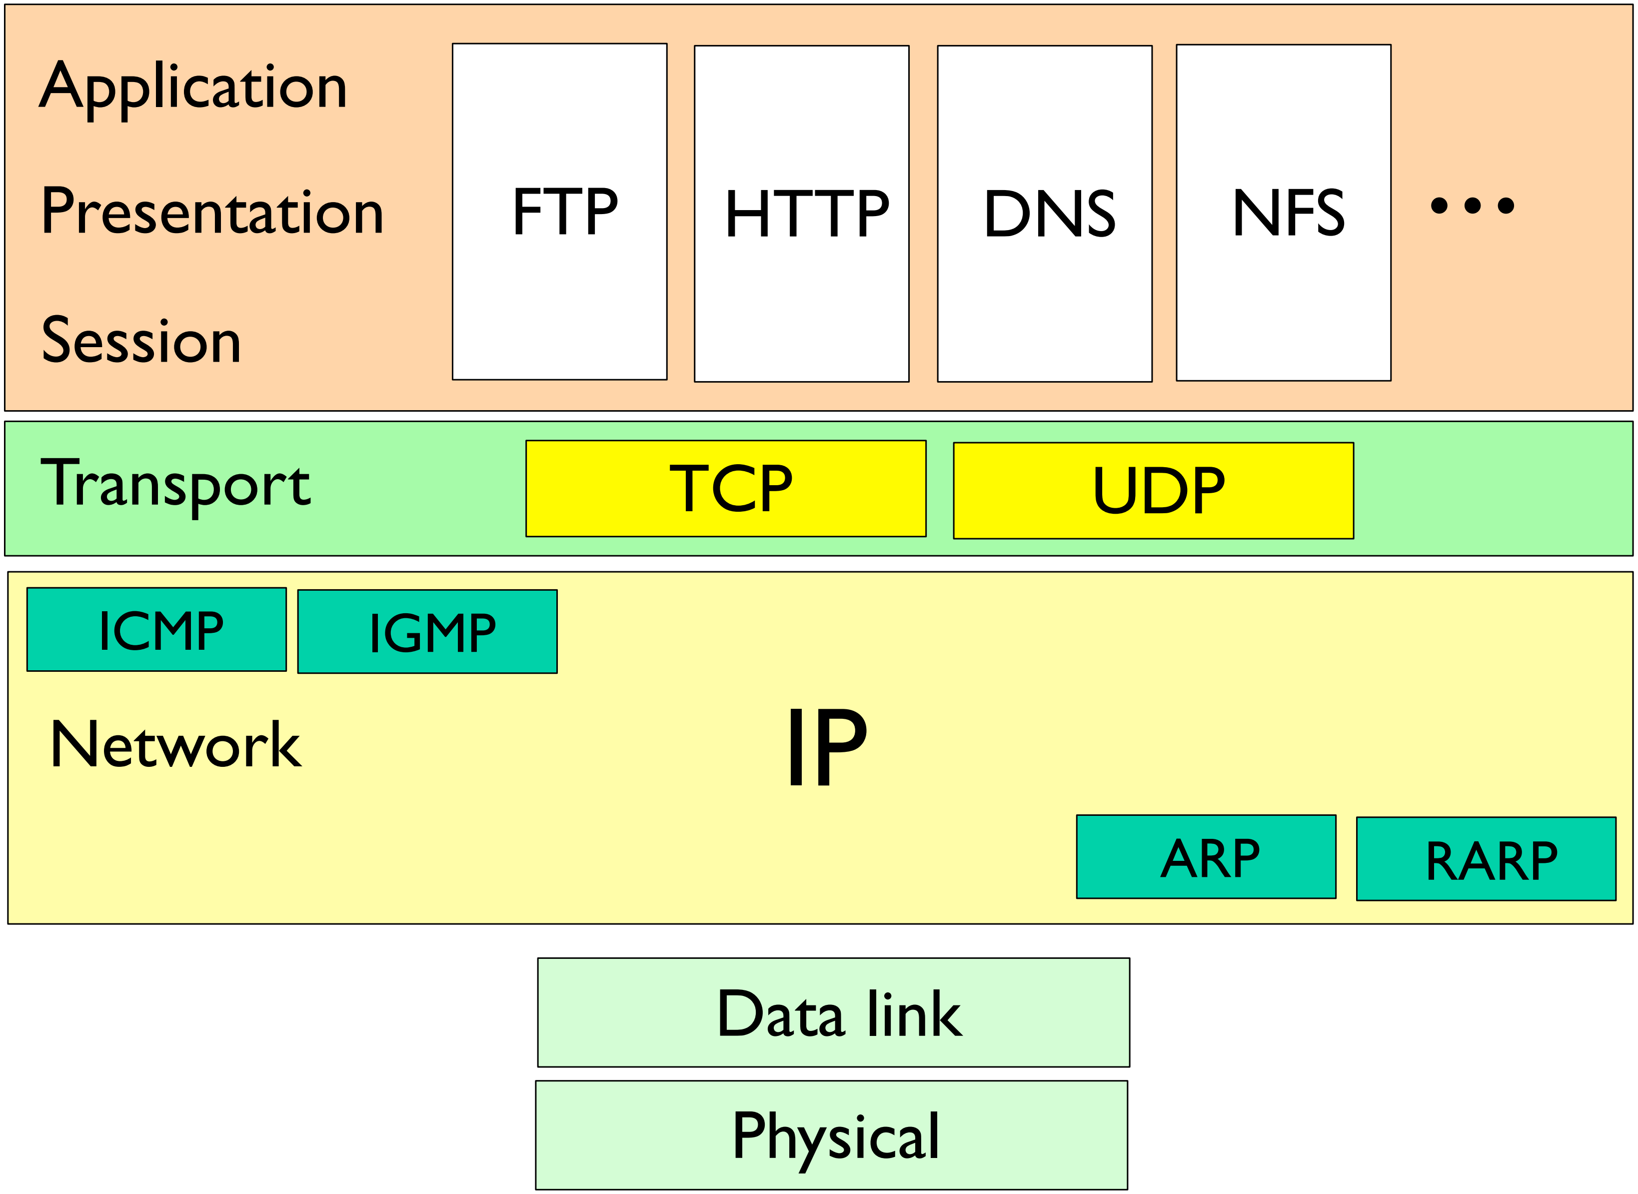

Performance

Bandwidth: number of bits per time unit.

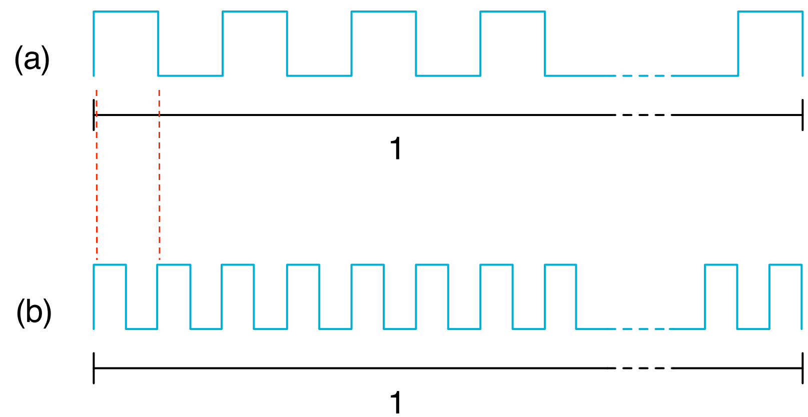

We can talk about bandwidth at the physical level, but we can also talk about logical process-to-process bandwidth.

Latency: time taken for a message to travel from one end of the network to the other.

Again, we can consider a single-link or an end-to-end channel.

Latency

$$\text{Latency} = \text{Propagation} + \text{Transmit} + \text{Queue}$$

$$\text{Propagation} = \text{Distance} / \text{Speed of light}$$

$$\text{Transmit} = \text{Size} / \text{Bandwidth}$$

$$\text{Speed of light} = \begin{cases} 2.0 \times 10^8 m/s \text{ in a fiber} \\ 2.3 \times 10^8 m/s \text{ in a cable} \\ 3.0 \times 10^8 m/s \text{ in a vaccum} \end{cases}$$

Delay x Bandwidth

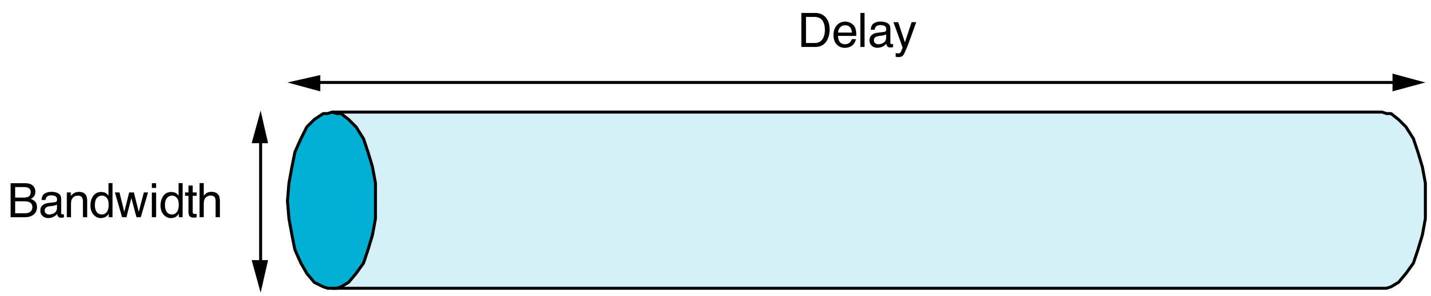

This product is analogous to the volume of a pipe or the number of bits it holds. It corresponds to how many bits the sender must transmit before the first bit arrives at the receiver.

Delay may be thought of as one-way latency or round-trip time (**RTT**) depending on the context.

Throughput

$$\text{Throughput} = \text{Transfer size} / \text{Transfer time}$$

(effective end-to-end throughput)

$$\text{Transfer time} = \text{RTT} + 1/\text{Bandwidth} \times \text{Transfer size}$$

Jitter

Jitter is a variation (somewhat random) of the latency from packet to packet. Jitter is most often observed when packets traverse multiple *hops* from source to destination.

Question: What is the cause of jitter?

Latency vs. RTT

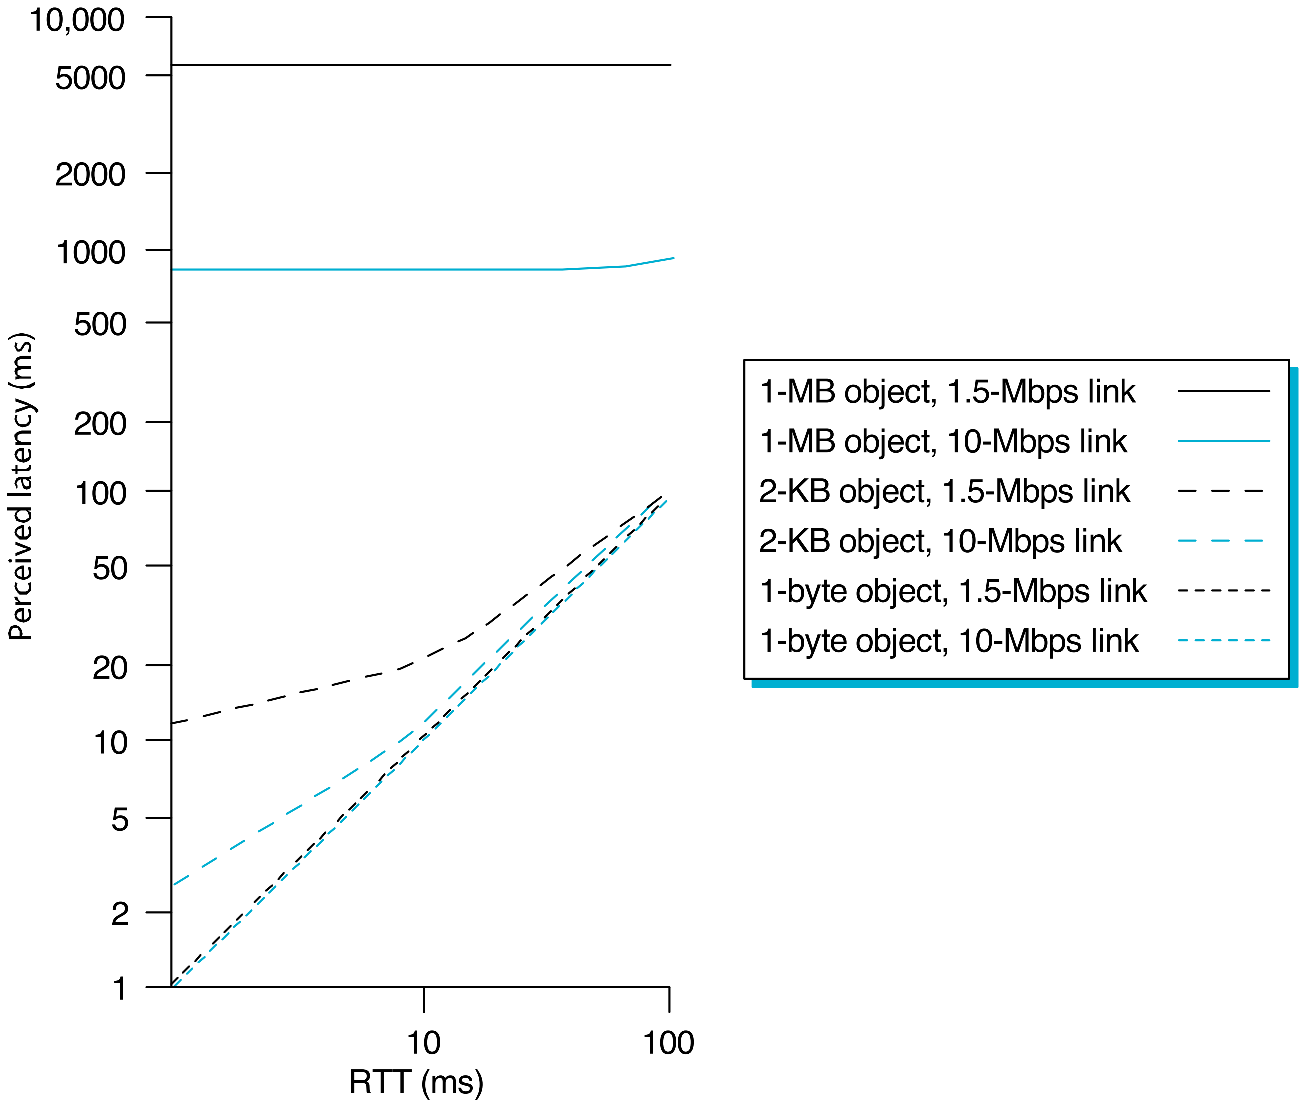

Building Blocks

- **Networks nodes**
- **Links**
  - –Dedicated cables
  - –Leased lines
  - –Last-mile links
  - –Wireless

Network Node

Memory: getting larger and larger, but never infinitely so.

Processor: Moore's law still holds for speed (memory latency improves much slower, however).

On a typical networked application, one must keep in mind the ***computation to communication ratio***.

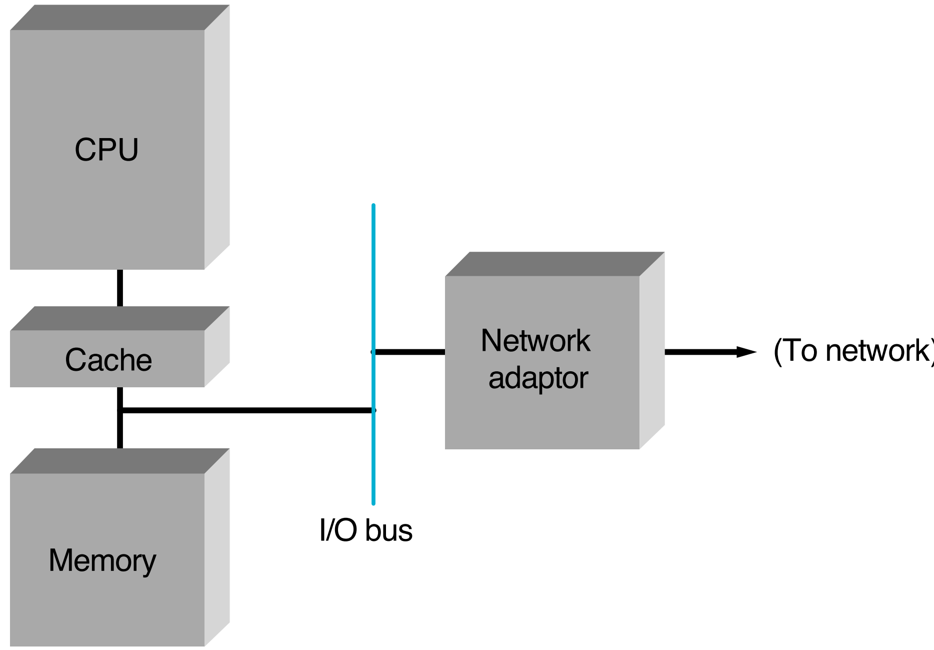

Links & Signals

Links: Twisted pair, coax, optical fiber, the *ether*; half-duplex or full-duplex.

Signals: Waveforms that travel on some medium (at the speed of light).

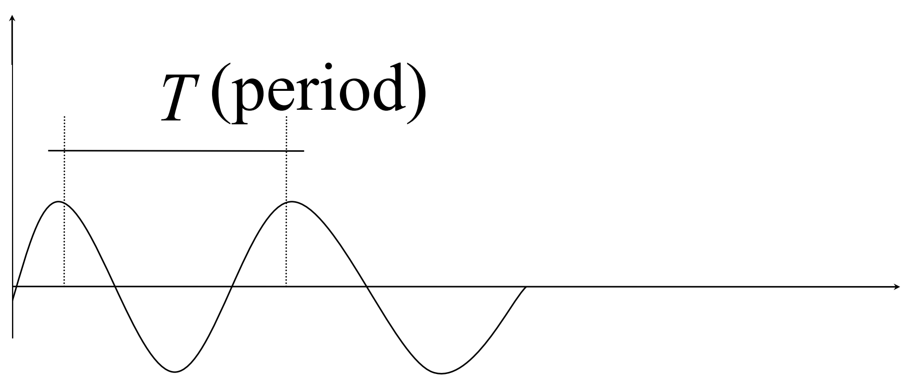

$$f = \frac{1}{T}$$

(frequency)

$$\lambda = c_{\text{medium}} f$$

(wavelength)

Spectrum

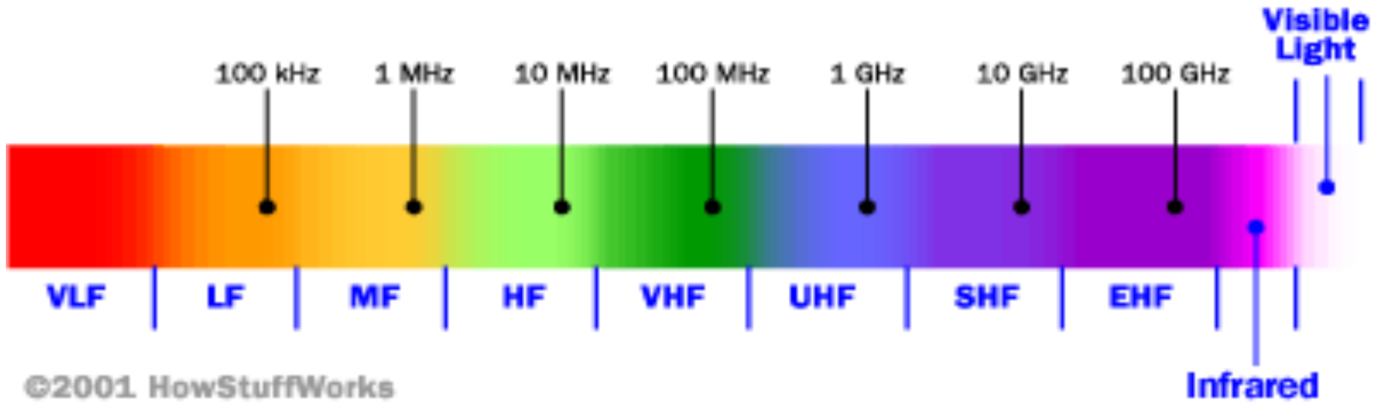

Visible Light
100 kHz
1 MHz
10 MHz
100 MHz
1 GHz
10 GHz
100 GHz
VLF
LF
MF
HF
VHF
UHF
SHF
EHF
Infrared
©2001 HowStuffWorks

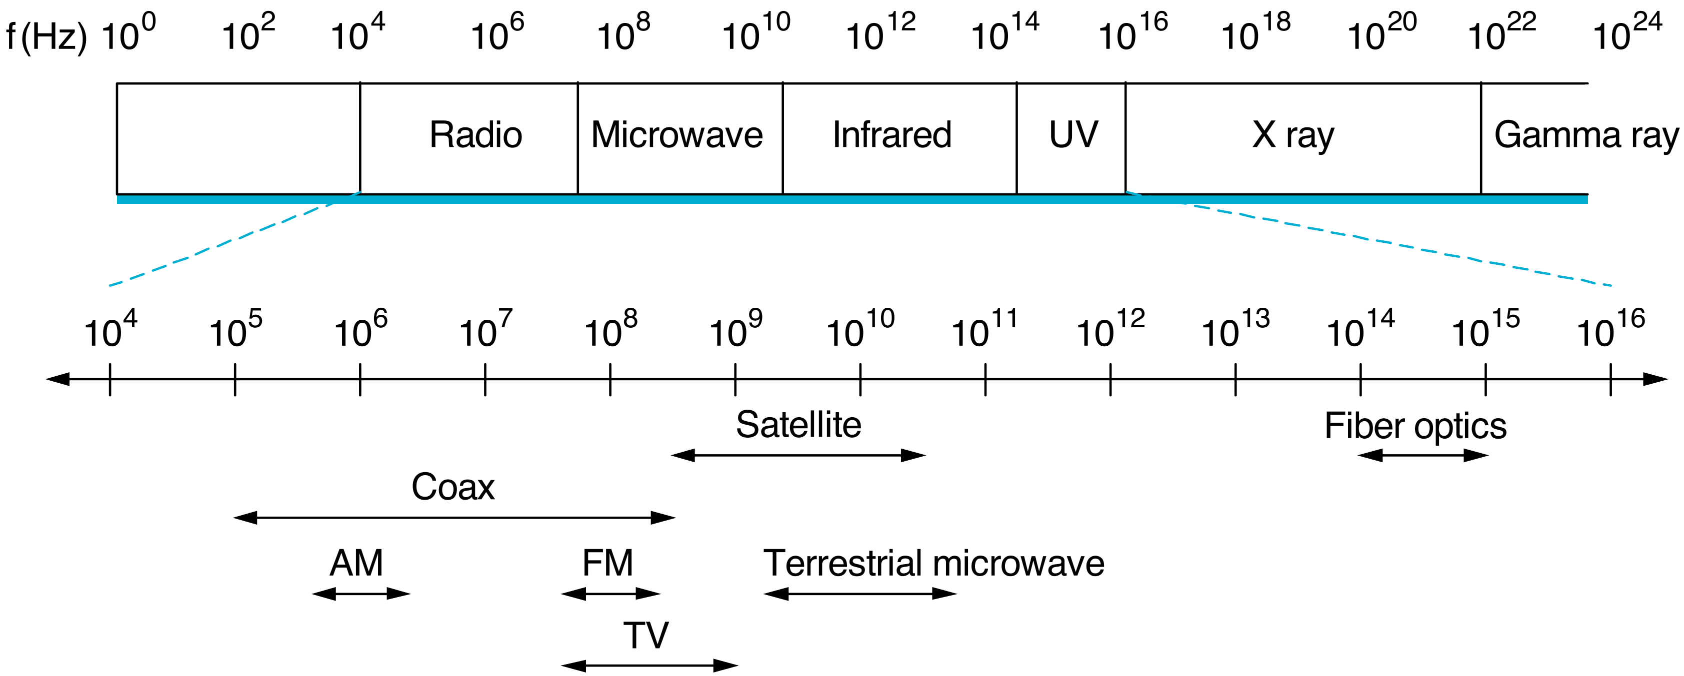

f (Hz) 10^0 10^2 10^4 10^6 10^8 10^10 10^12 10^14 10^16 10^18 10^20 10^22 10^24
Radio
Microwave
Infrared
UV
X ray
Gamma ray
10^4 10^5 10^6 10^7 10^8 10^9 10^10 10^11 10^12 10^13 10^14 10^15 10^16
Satellite
Fiber optics
Coax
AM
FM
Terrestrial microwave
TV