Computer Networking

Introduction

Quintin jean-noël

Jean-noel.quintin@imag.fr

Grenoble university

Based on the presentation of Duda

http://duda.imag.fr

Course organization

- Introduction
  - Network and architecture
  - OSI model
  - Performances
- Application Layer
  - HTTP, SMTP, DNS protocols
- Transport Layer
  - Notion de socket
  - UDP, TCP
- TCP
  - flux, congestion
- Network Layer
  - IP
  - route Tables
- Physical Link Layer

Course support

- Web site
  - **`http://duda.imag.fr/2at`**
- J. Kurose, K. Ross “Computer Networking”, 4th edition, Addison Wesley, 2007
- J. Kurose, K. Ross, "Analyse structurée des réseaux. Des applications de l'internet aux infrastructures des télécommunications." Pearson Education France, 2003
- Others
  - L. Toutain "Réseaux locaux et Internet", 3me édition, Hermes, 2003
  - W. R. Stevens “TCP/IP illustrated, Volume I”, Addison Wesley (Very detailed, experimental hands-on description of TCP/IP)

Overview

- Network architectures
  - recall on the Internet
  - protocol architectures
    - how entities cooperate?
  - interconnection structure
    - which entities are connected?
  - related protocols
    - how and where different functionalities are implemented?
- Performance
  - transmission
  - propagation
  - bandwidth-delay product
  - queueing delay

Inside the Internet

- Between end systems
  - TCP protocol for reliable transmission
- Inside the network core
  - IP protocol: forwarding packets between routers
- Between routers or between end system and router
  - high speed link: ATM, POS (Packet over SONET), satellite links
  - access network: Ethernet, modem, xDSL, HFC

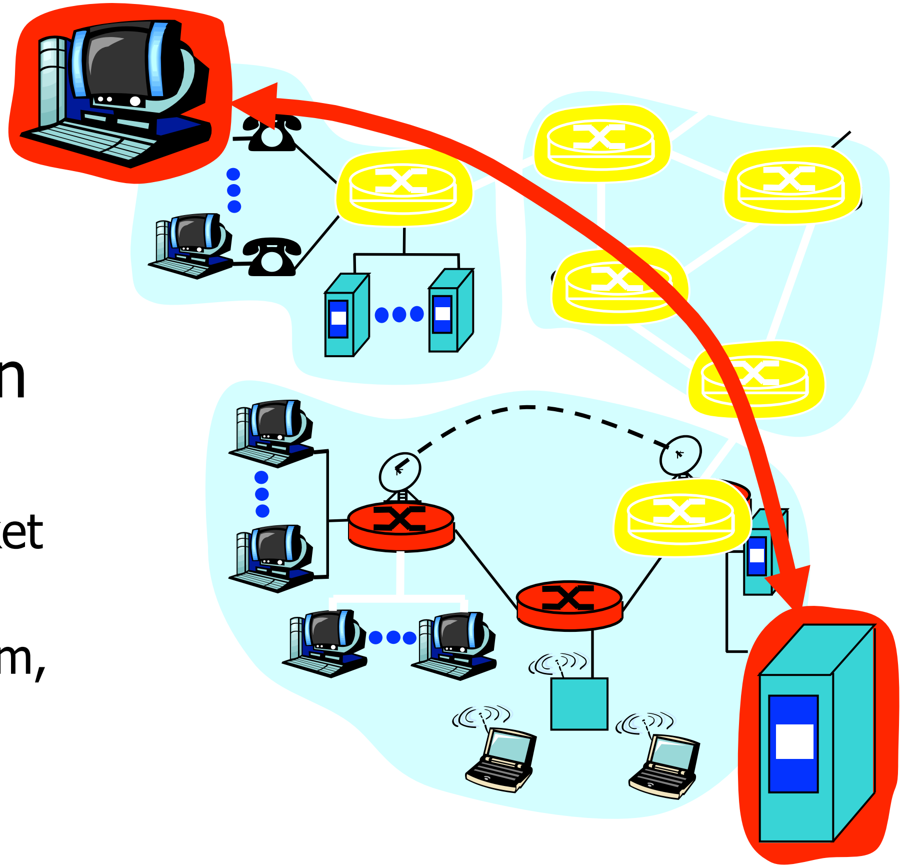

Network structure

- network edge: applications and hosts
- network core:
  - routers
  - network of networks
- access networks, physical media: communication links

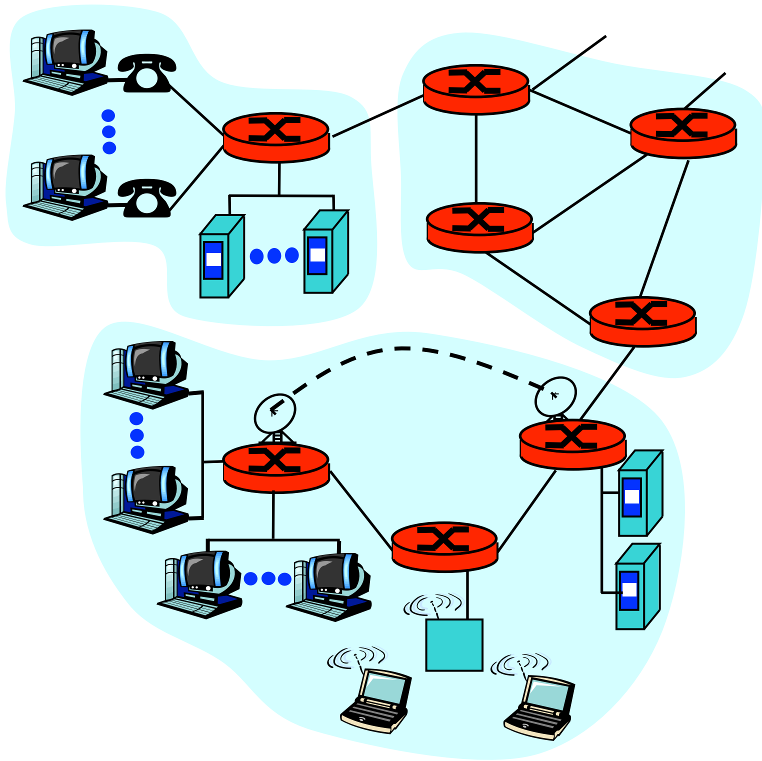

The network edge:

- end systems (hosts):
  - run application programs
  - e.g., WWW, email
  - at "edge of network"
- client/server model
  - client host requests, receives service from server
  - e.g., WWW client (browser)/ server, email client/server
- peer-peer model:
  - symmetric host interaction
  - e.g. teleconferencing

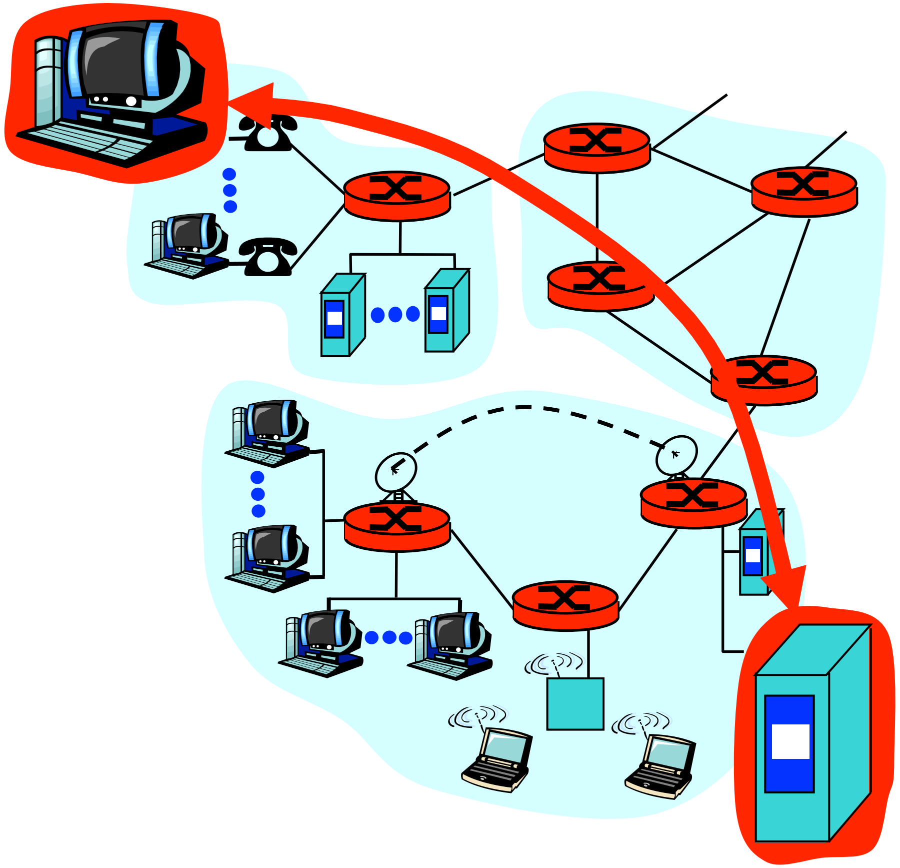

The Network Core

- mesh of interconnected routers
- *the* fundamental question: how is data transferred through net?
  - circuit switching: dedicated circuit per call: telephone nets
  - packet-switching: data sent thru net in discrete "chunks" (IP)

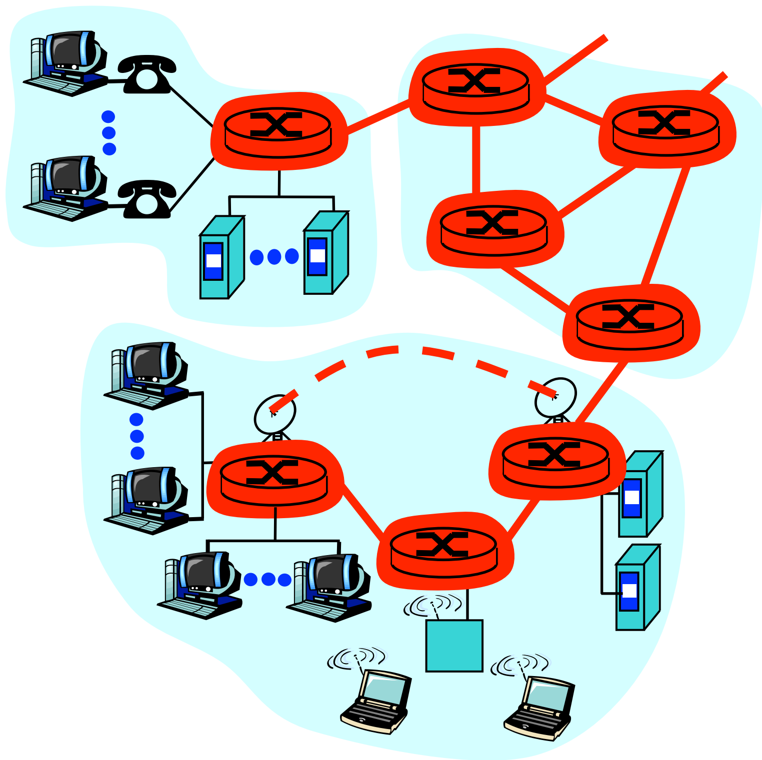

Access networks and physical media

How to connect end systems to edge router?

- residential access nets
- institutional access networks (school, company)
- mobile access networks

Characteristics:

- bandwidth (bits per second) of access network
- shared or dedicated

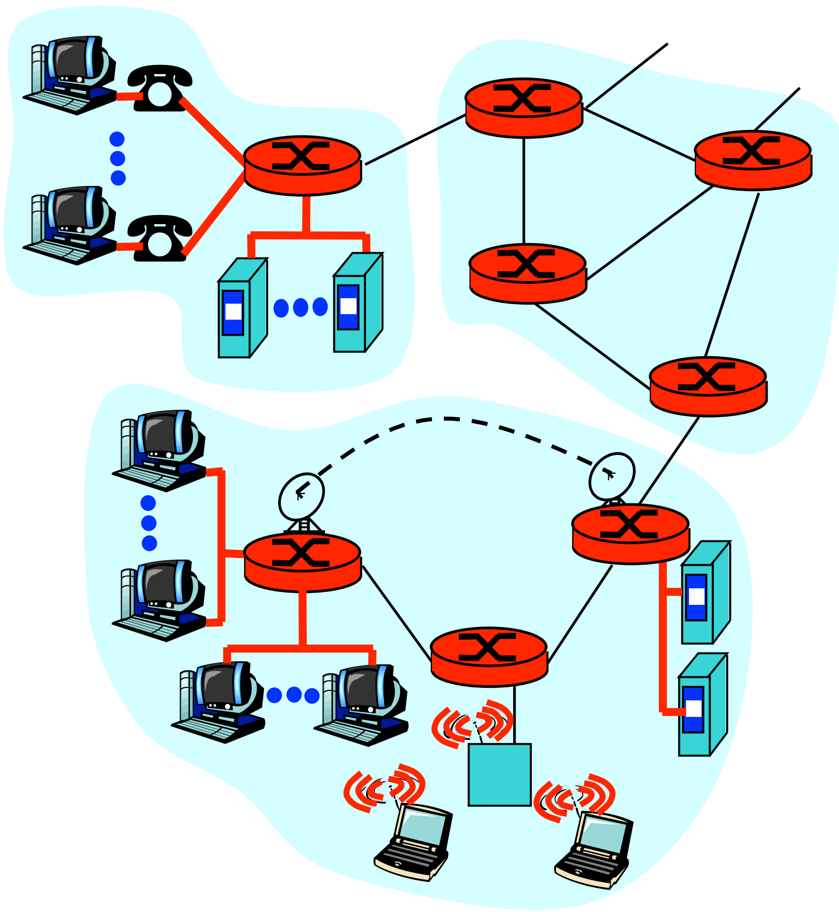

How is a communication done?

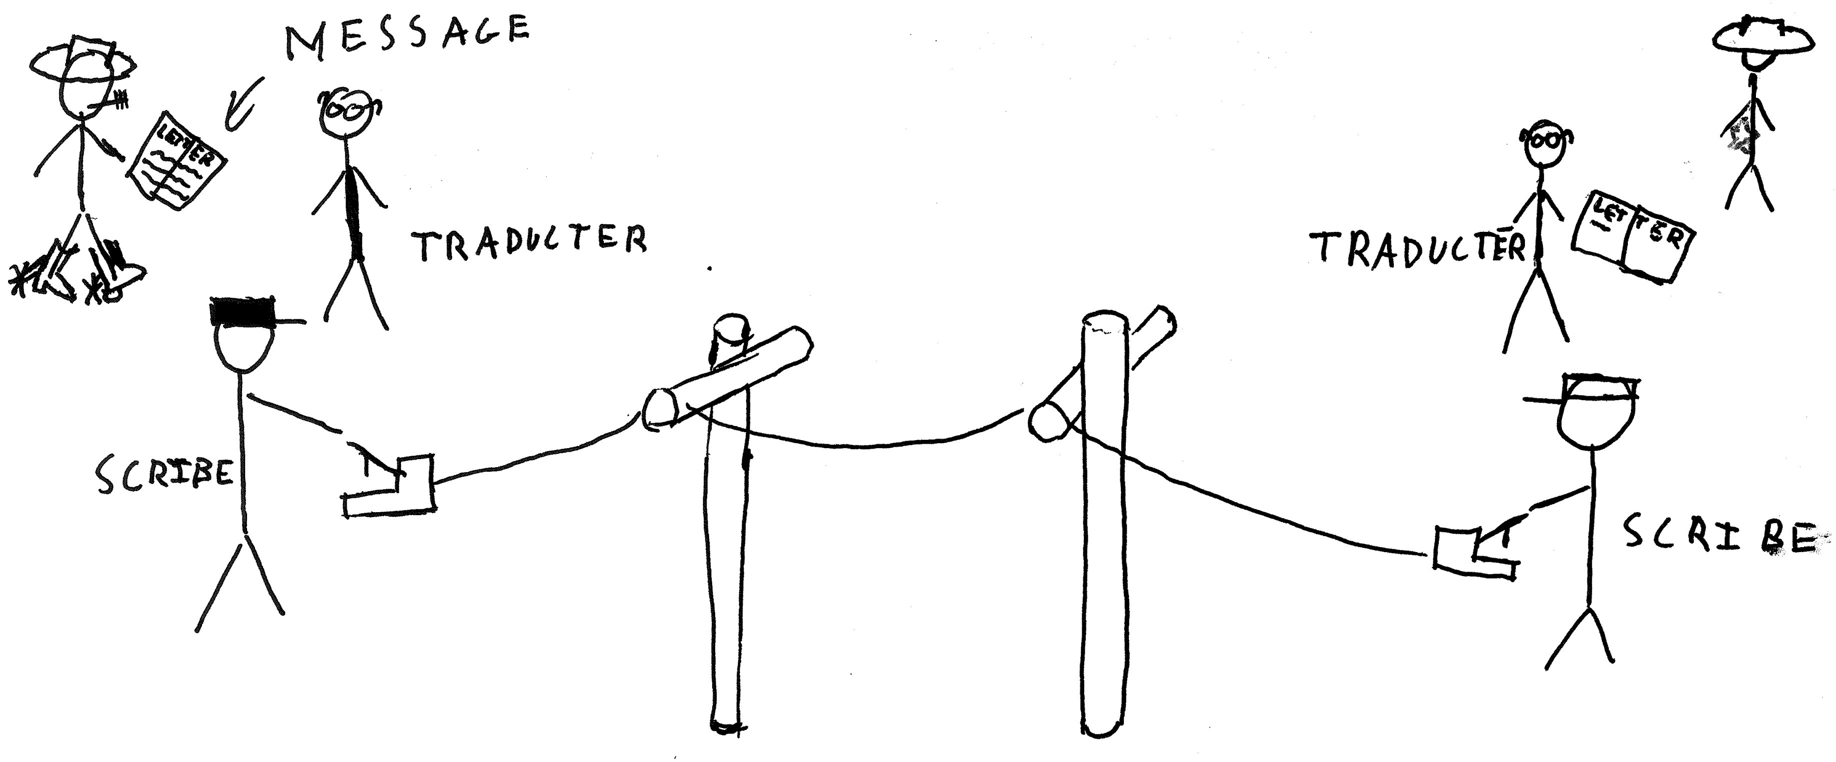

TCP/IP Architecture

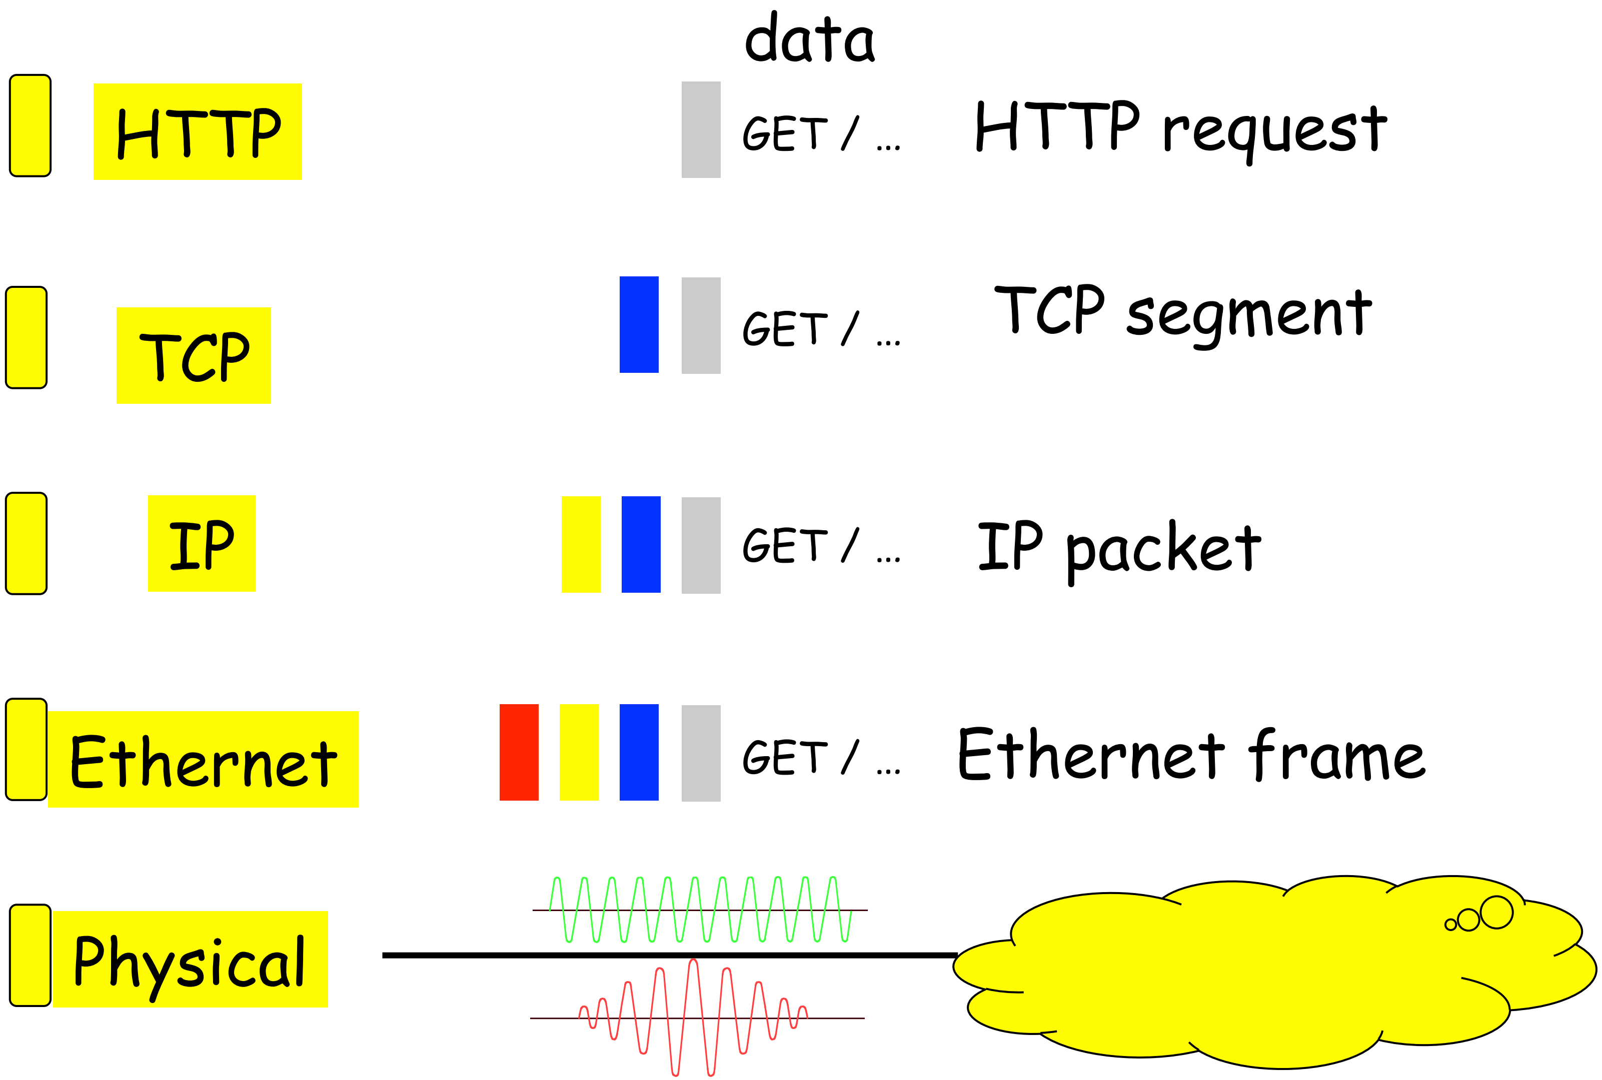

Application Layer

- Application layer supports network application
  - applications that are distributed over the network
  - applications that communicates through the network
- Many known protocols
  - FTP: file transfer
  - SMTP: email protocol
  - HTTP:web protocol
- An application uses UDP or TCP, it is a designer’s choice
- Interface with the transport layer
  - use for example the **socket** API: a library of C functions
  - **socket** also means (IP address, port number)

Transport Layer

- Why a transport layer ?
  - **transport layer** = makes network service available to programs
  - is end-to-end only, not in routers
- In TCP/IP there are two transport protocols
  - UDP (user datagram protocol)
    - unreliable
    - offers a datagram service to the application (unit of information is a message)
  - TCP (transmission control protocol)
    - reliable
    - offers a stream service (unit of information is a byte)

Layering: *logical* communication

- E.g.: transport
- take data from app
- add addressing, reliability check info to form "datagram"
- send datagram to peer
- wait for peer to ack receipt
- analogy: post office

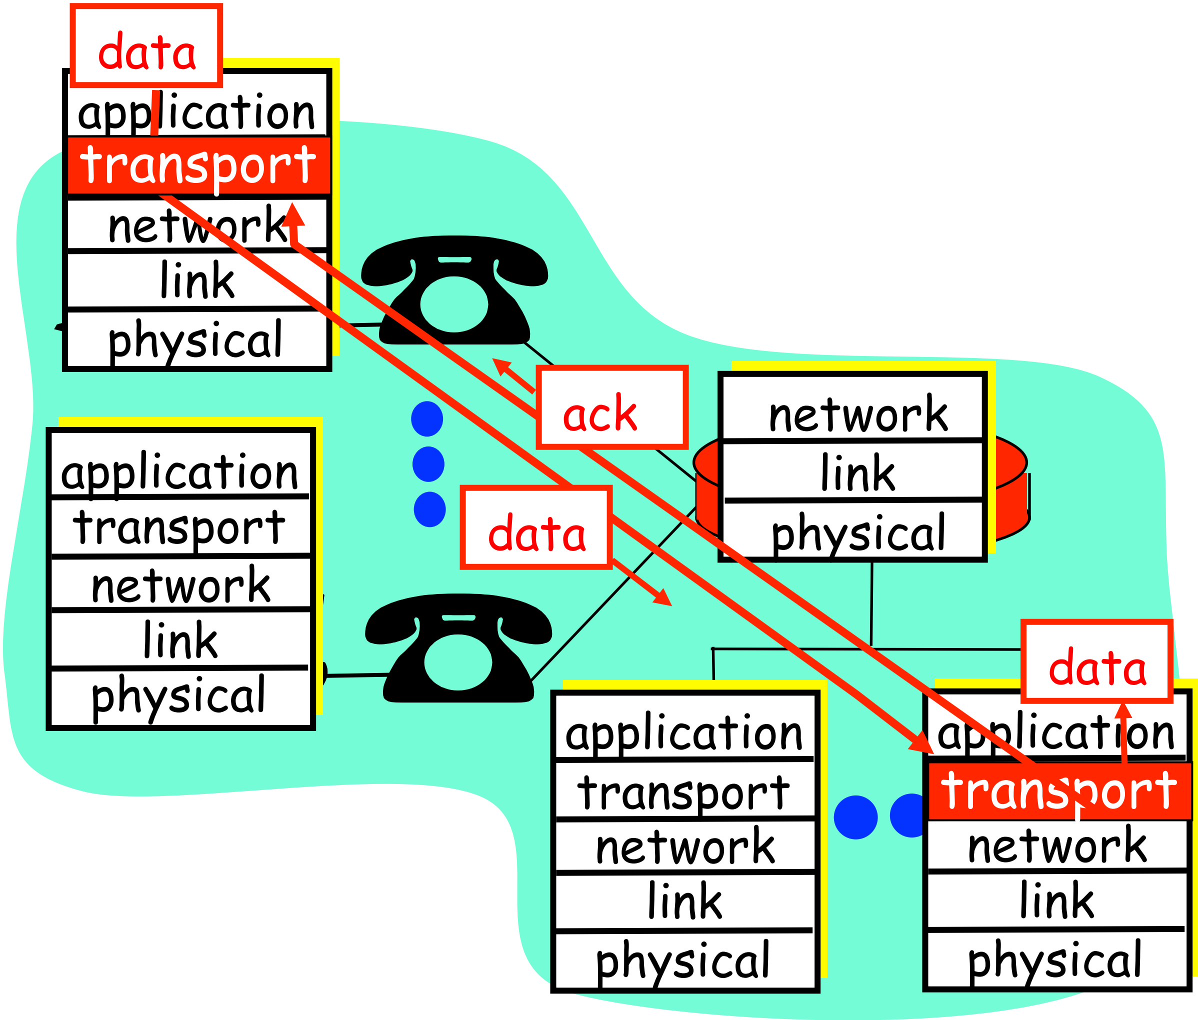

TCP

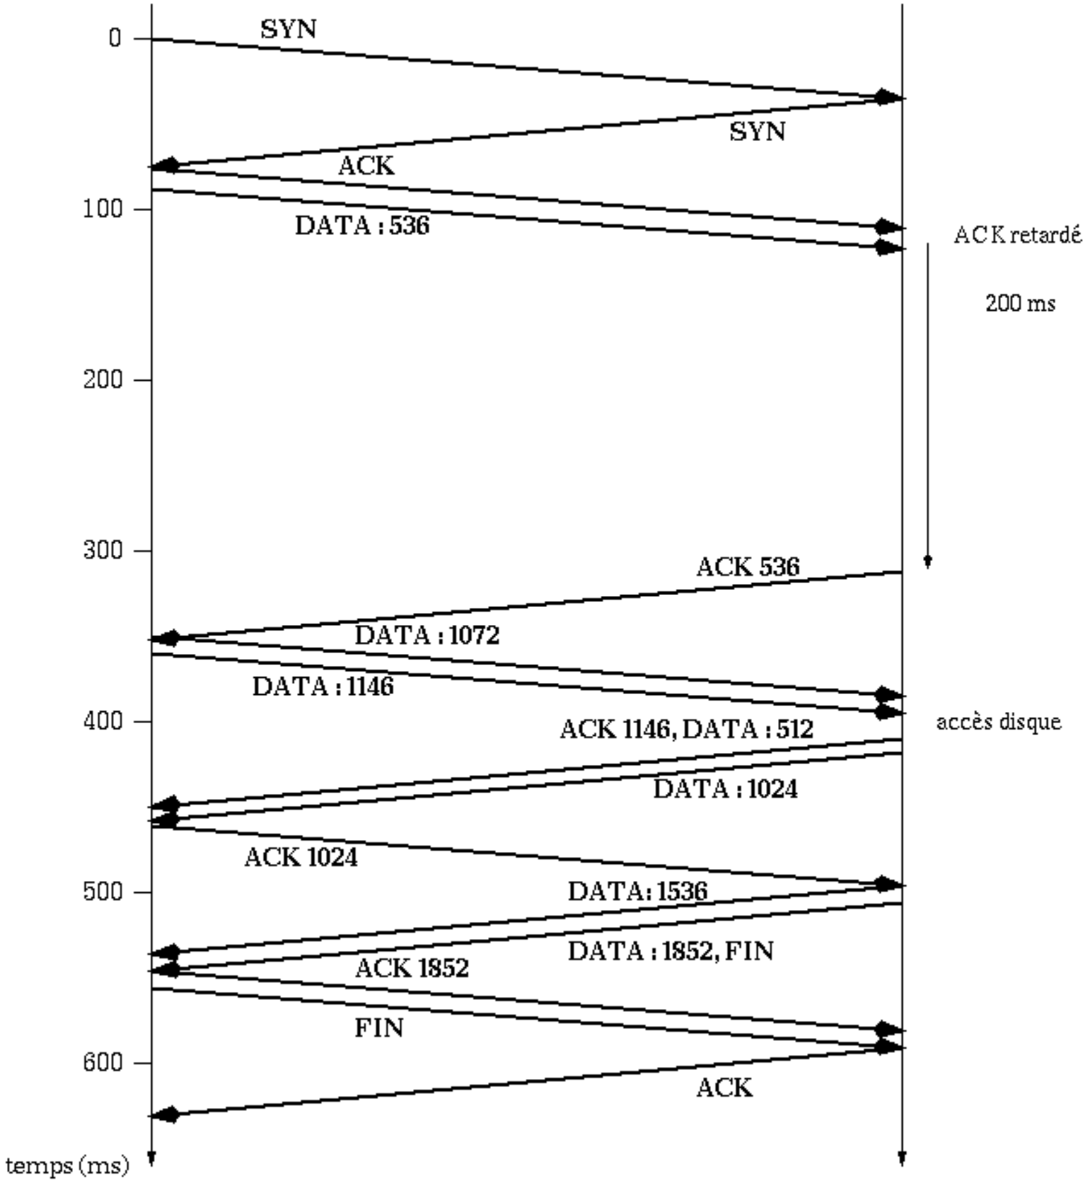

HTTP Request : 1146 octets

HTTP Reply : 1852 octets

Taille maximale de segment : 536 octets

Temps aller-retour (RTT) : ~70 ms
(North Carolina -> Illinois)

Problèmes :

- établissement de connexion
- démarrage lent (Slow Start)
- ACK retardé
- temps aller-retour

Network Layer

- Set of functions required to transfer packets end-to-end (from host to host)
  - hosts are not directly connected - need for intermediate systems
  - examples: IP, Appletalk, IPX
- Intermediate systems
  - routers: forward packets to the final destination
  - interconnection devices

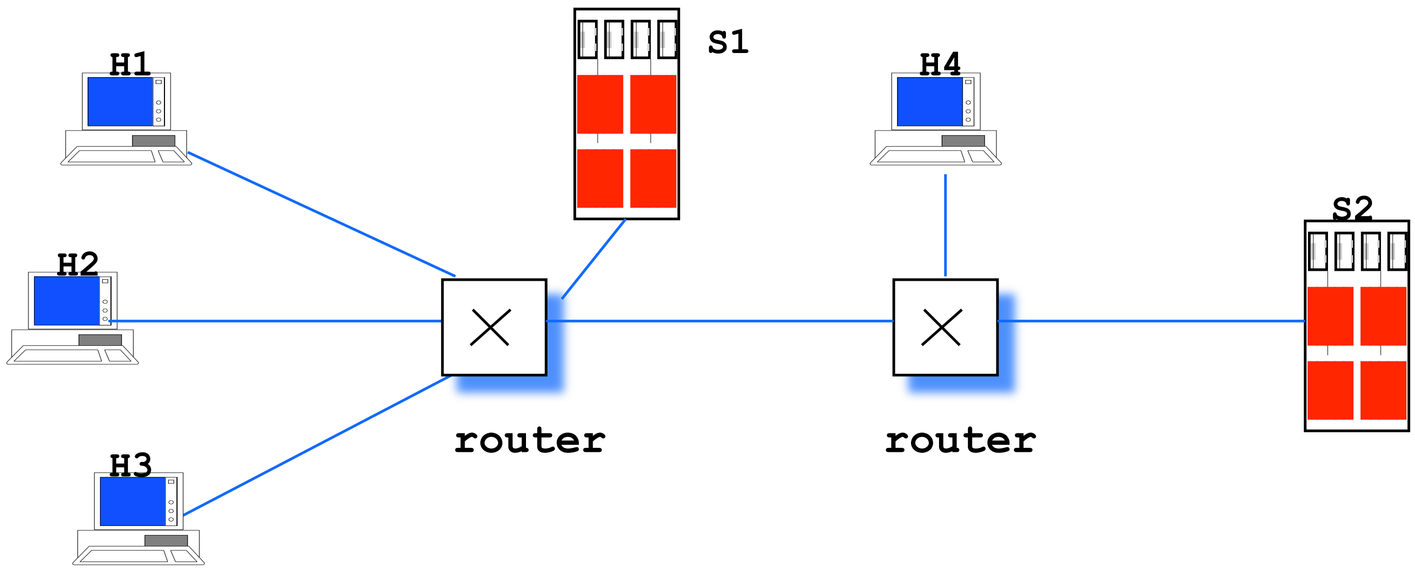

IP

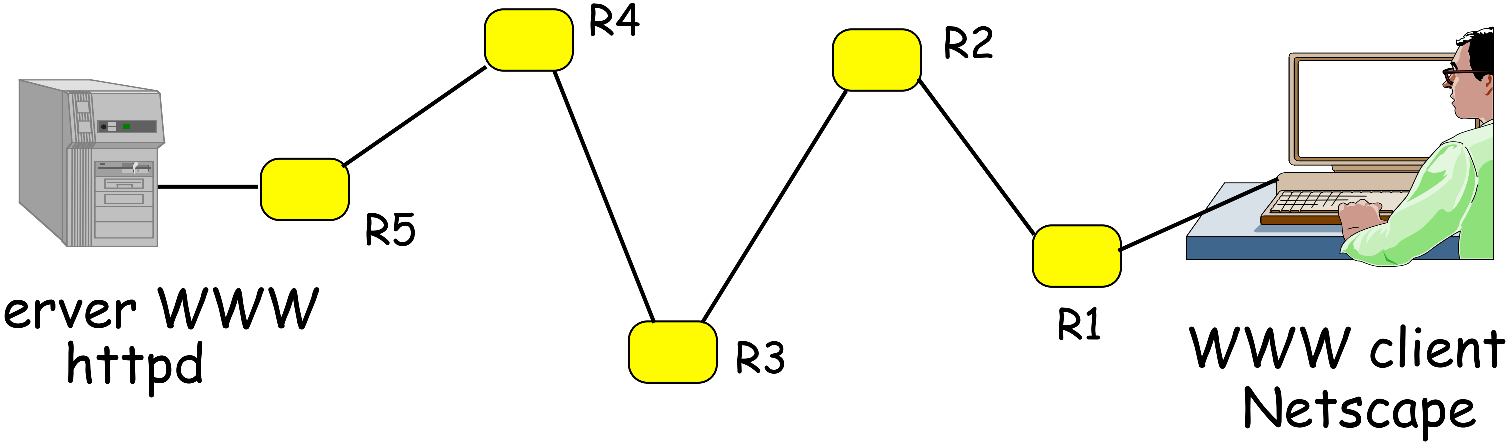

Routing Table of R3

dest	router
129.88.38	R4

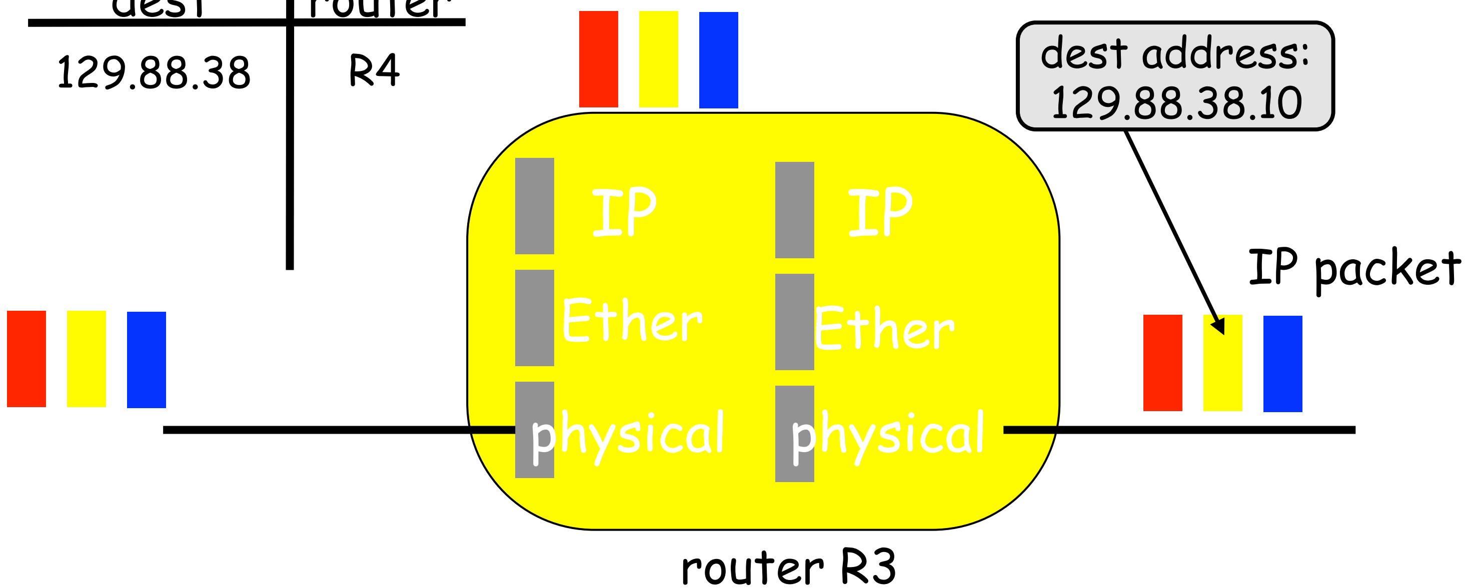

Physical Layer
Data Link Layer

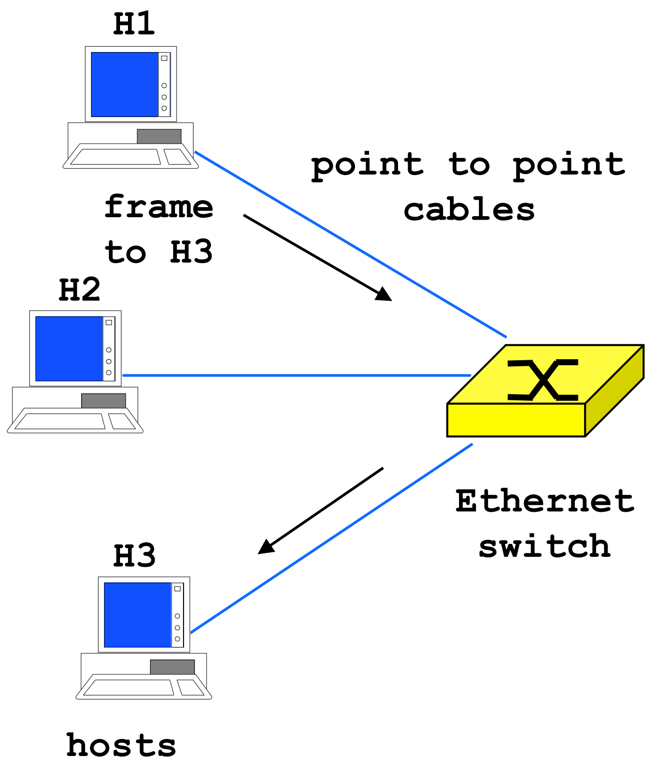

- Physical transmission = **Physical** function
  - bits <-> electrical / optical signals
  - transmit individual bits over the cable: modulation, encoding
- Frame transmission = **Data Link** function
  - bits <-> frames
  - bit error detection
  - packet boundaries
  - in some cases: error correction by retransmission (802.11)
- Modems, xDSL, LANs

Protocol architectures

- Protocol entity
  - provides a set of services, e.g..
    - *connect, send*
  - data multiplexing/demultiplexing
  - construction/analysis of PDUs
  - execution of procedures
- Protocol unit (PDU)
  - header: control functions
  - opaque data
- Procedures
  - actions to perform protocol functions: e.g. lost packet retransmission

Internet protocol stack

- **Application**: supporting network applications
  - FTP, SMTP, HTTP, OSPF, RIP
- **Transport**: host-host data transfer
  - TCP, UDP
- **Network**: routing of datagrams from source to destination
  - IP
- **Link**: data transfer between neighboring network elements
  - PPP, Ethernet
- **Physical**: bits “on the wire”

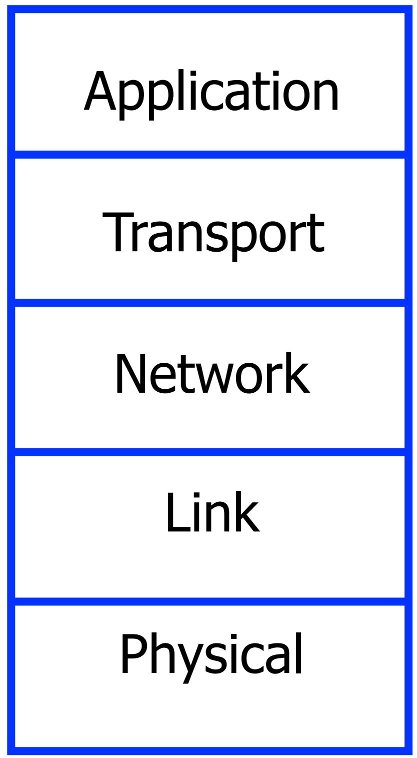

Layering: physical communication

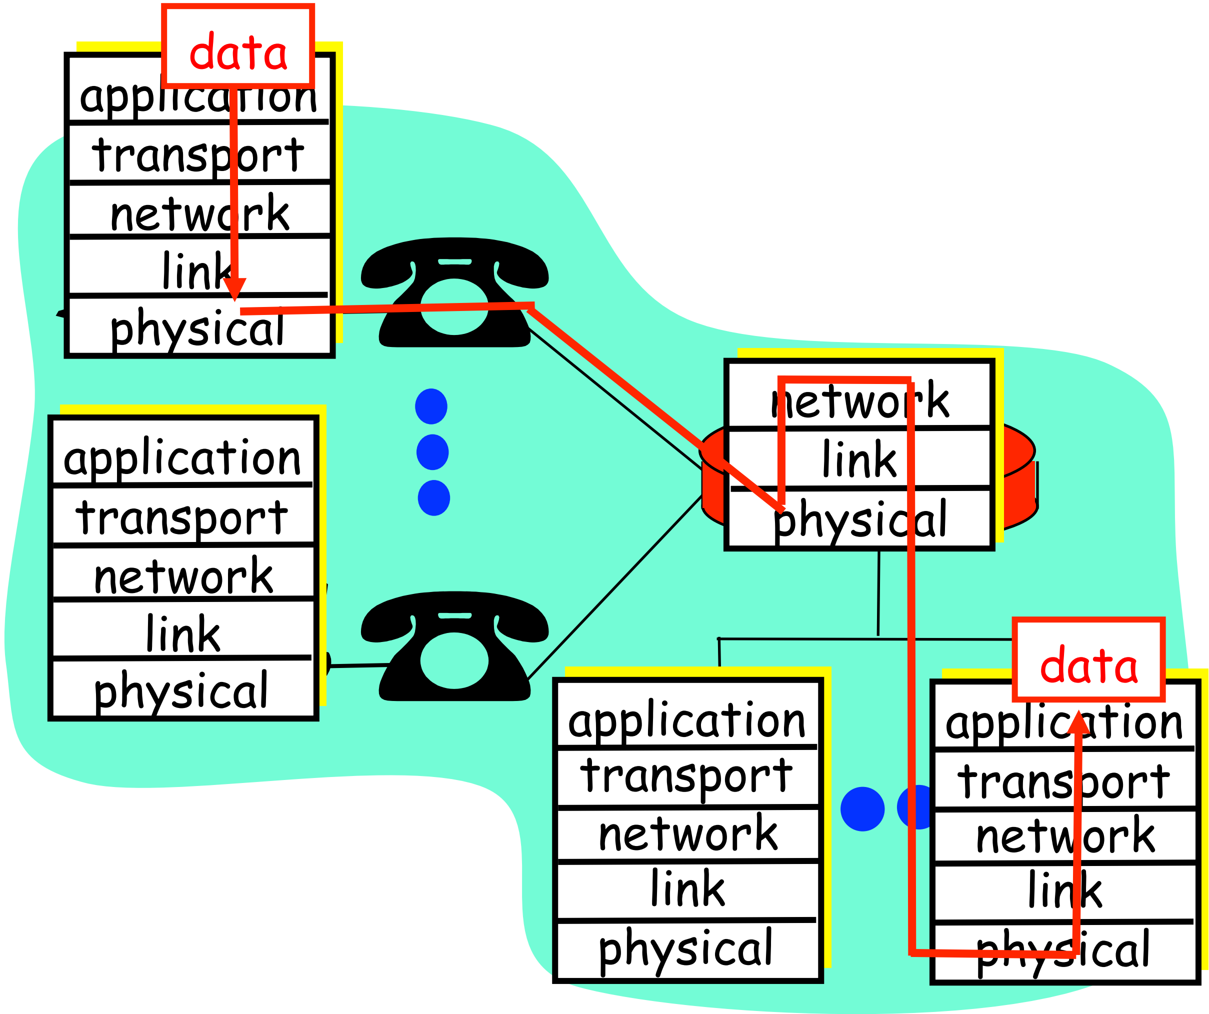

Protocol layering and data

Each layer takes data from above

- adds header information to create new data unit
- passes new data unit to layer below

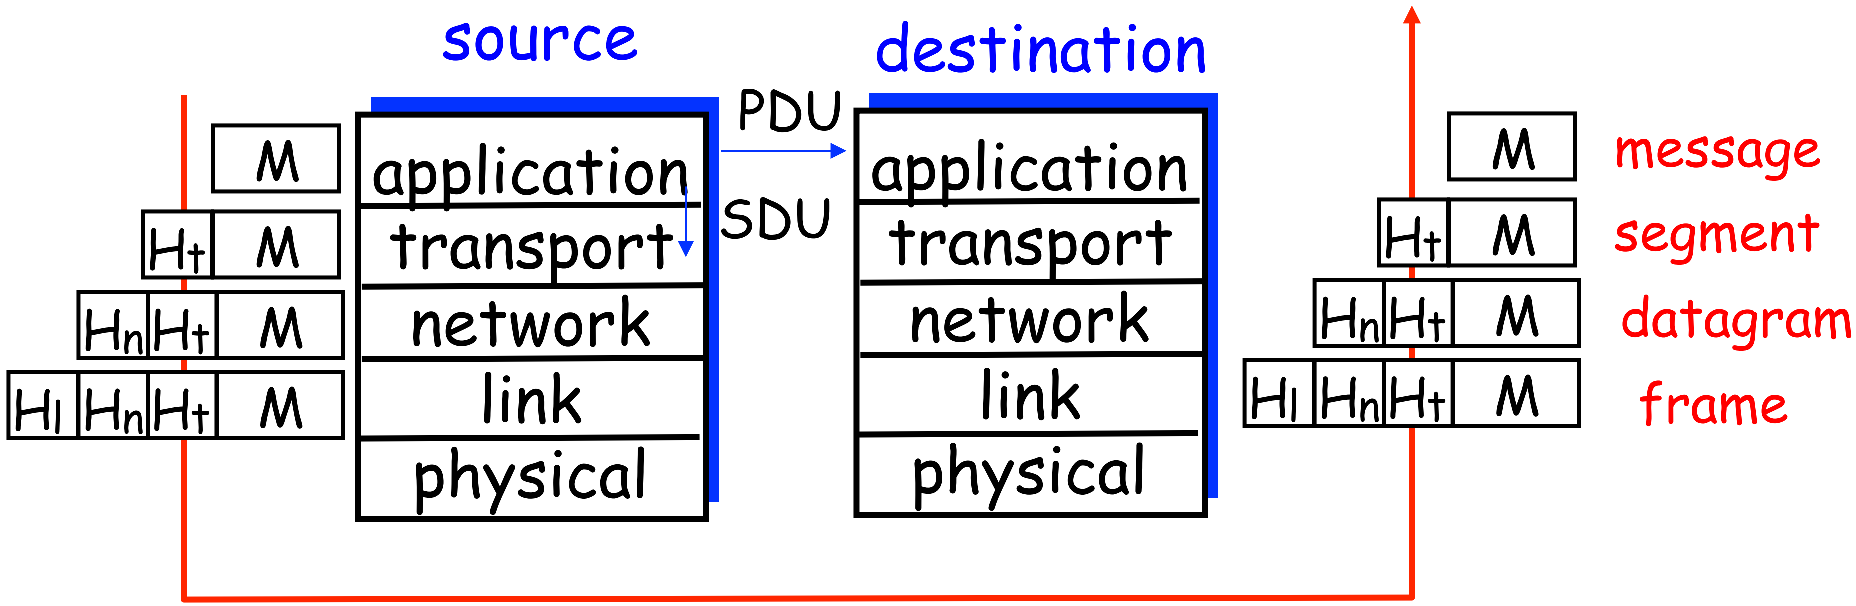

Encapsulation

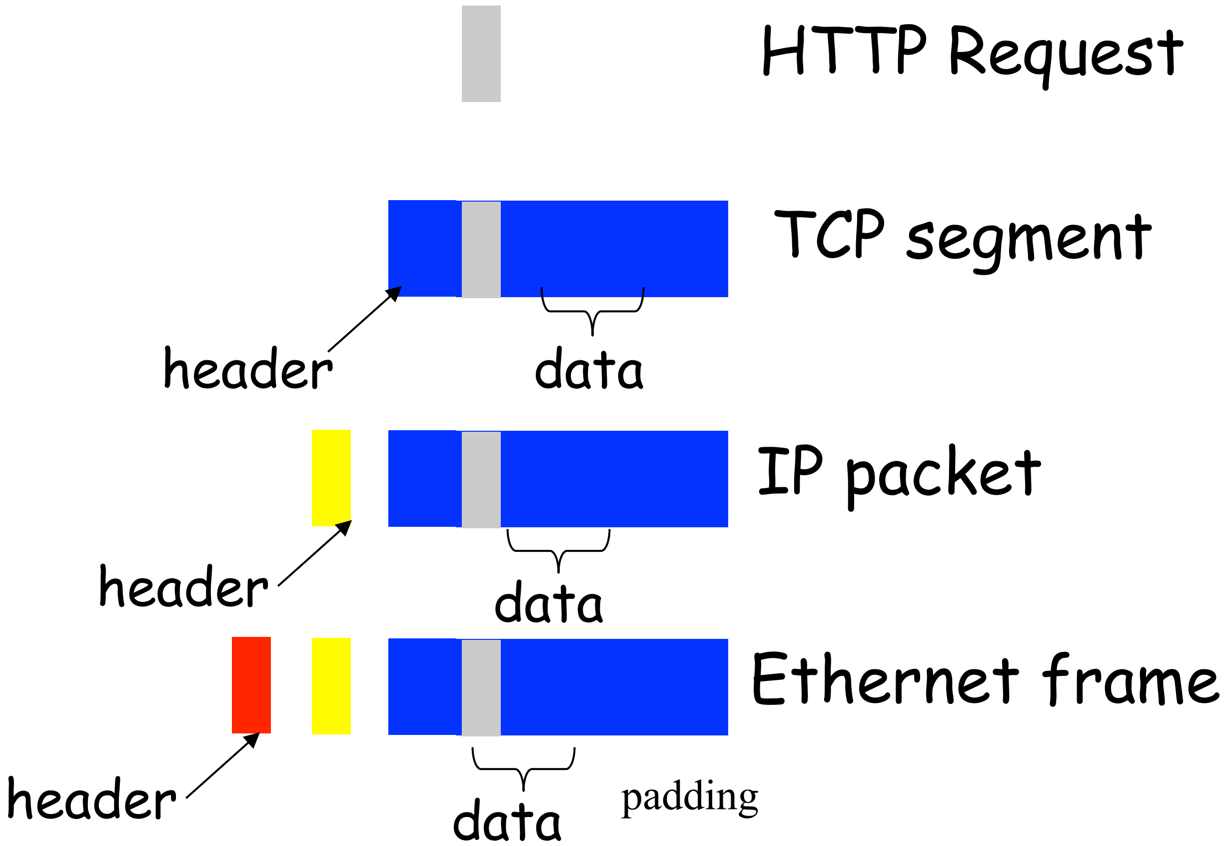

Ethernet Frame

- Length of data in the data link Layer
  - Minimal size is 46 bytes
    - The Paddig is a feild of insignificant data to obtain the minimal size
  - Maximal size is 1500 bytes
    - MTU (Maximal transmission unit):
      is the name for the maximal size for each Layer

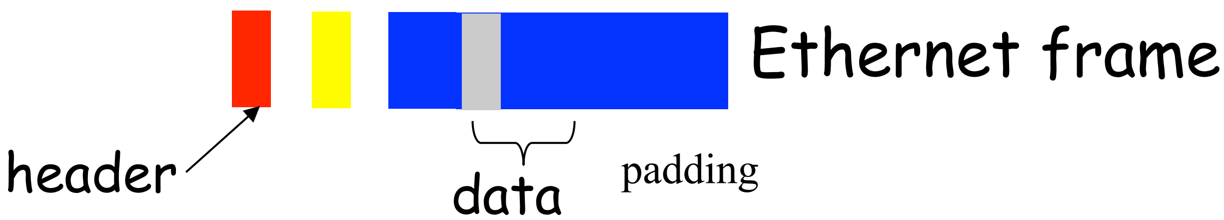

OSI ISO Model

Application	Common functions
Presentation	Interchangeable formats
Session	Organizing dialog
Transport	Reliable transmission
Network	Forwarding in the network
Data link	Transmission between two nodes
Physical	Signal transmission

QoS (Quality of Service)

- Different applications need different level of QoS

Application	Requires	Tolerates
e-mail	reliability	delay
file transfer	reliability	latency
VoIP	delay and jitter	loss
Streaming video	delay and bandwidth	loss
telesurgery	delay and bandwidth	jitter
instant messenging	delay	low bandwidth

LAN stack

- **Management**: e.g. construct forwarding tables
  - SNAP: Spanning Tree protocol
- **LLC**: multiplex different protocols
  - IP, IPX, SNAP
- **MAC**: medium access
  - 802.3 (Ethernet), 802.4 (Token Ring), 802.5 (Token Bus), 802.11 (Wi-Fi)
- **Physical**: bits "on the wire"

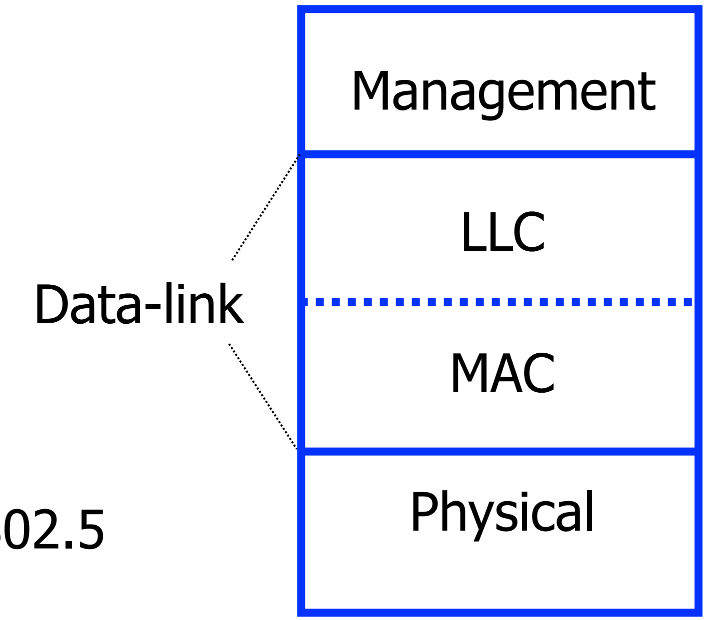

Interconnection structure - layer 3

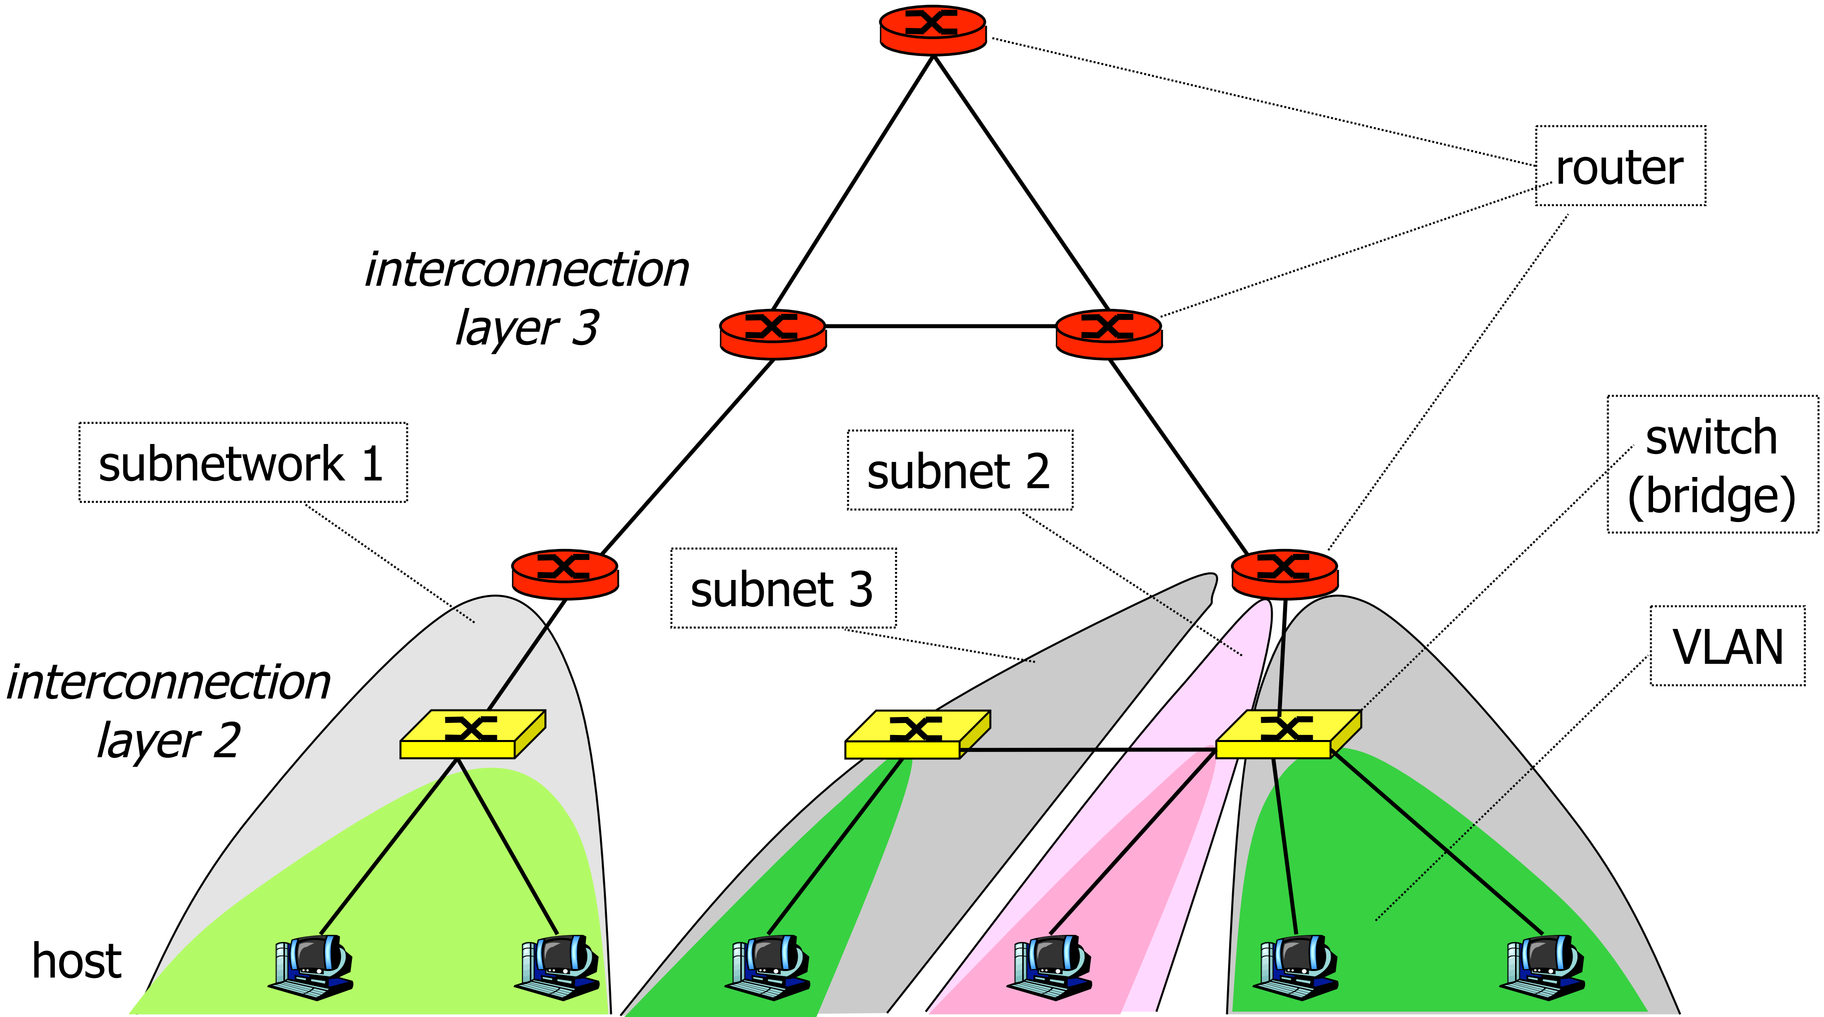

Interconnection at layer 2

- Switches (bridges)
  - interconnect hosts
  - logically separate groups of hosts (VLANs)
  - managed by one entity
- Type of the network
  - broadcast
- Forwarding based on MAC address
  - flat address space
  - forwarding tables: one entry per host
  - works if no loops
    - careful management
    - Spanning Tree protocol
  - not scalable

Protocol architecture

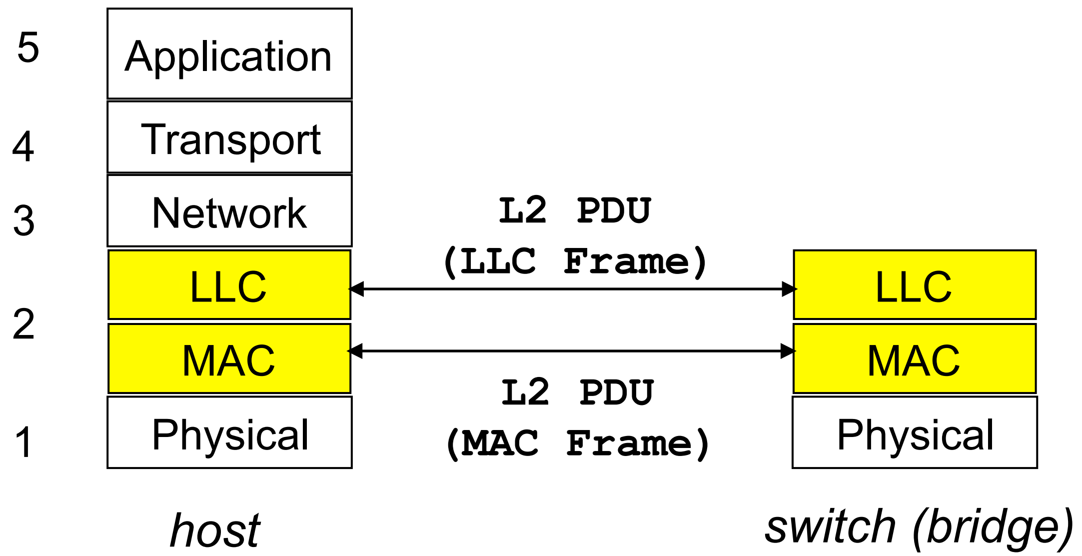

- Switches are layer 2 intermediate systems
- Transparent forwarding
- Management protocols (Spanning Tree, VLAN)

Protocols

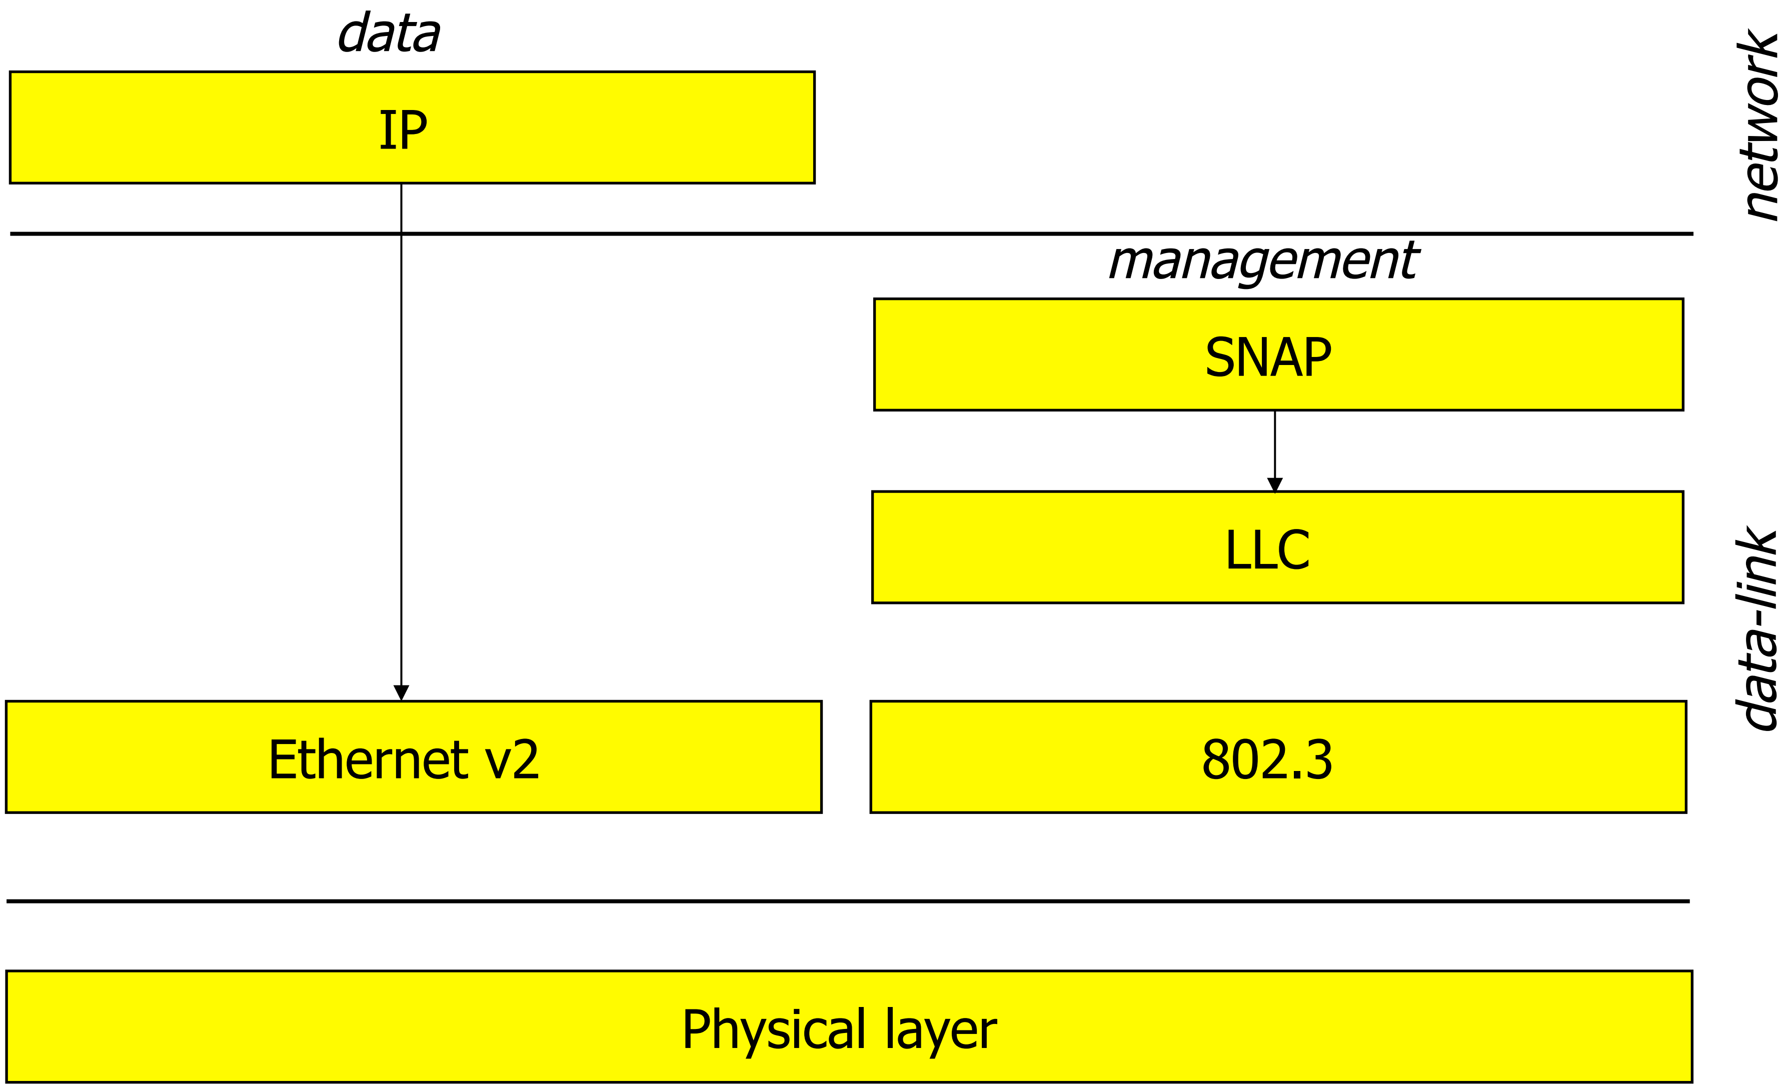

data
IP
network
management
SNAP
LLC
data-link
Ethernet v2
802.3
Physical layer

Interconnection at layer 3

- Routers
  - interconnect subnetworks
  - logically separate groups of hosts
  - managed by one entity
- Forwarding based on IP address
  - structured address space
  - routing tables: aggregation of entries
  - works if no loops - routing protocols (IGP - Internal Routing Protocols)
  - scalable inside one administrative domain

Protocol architecture

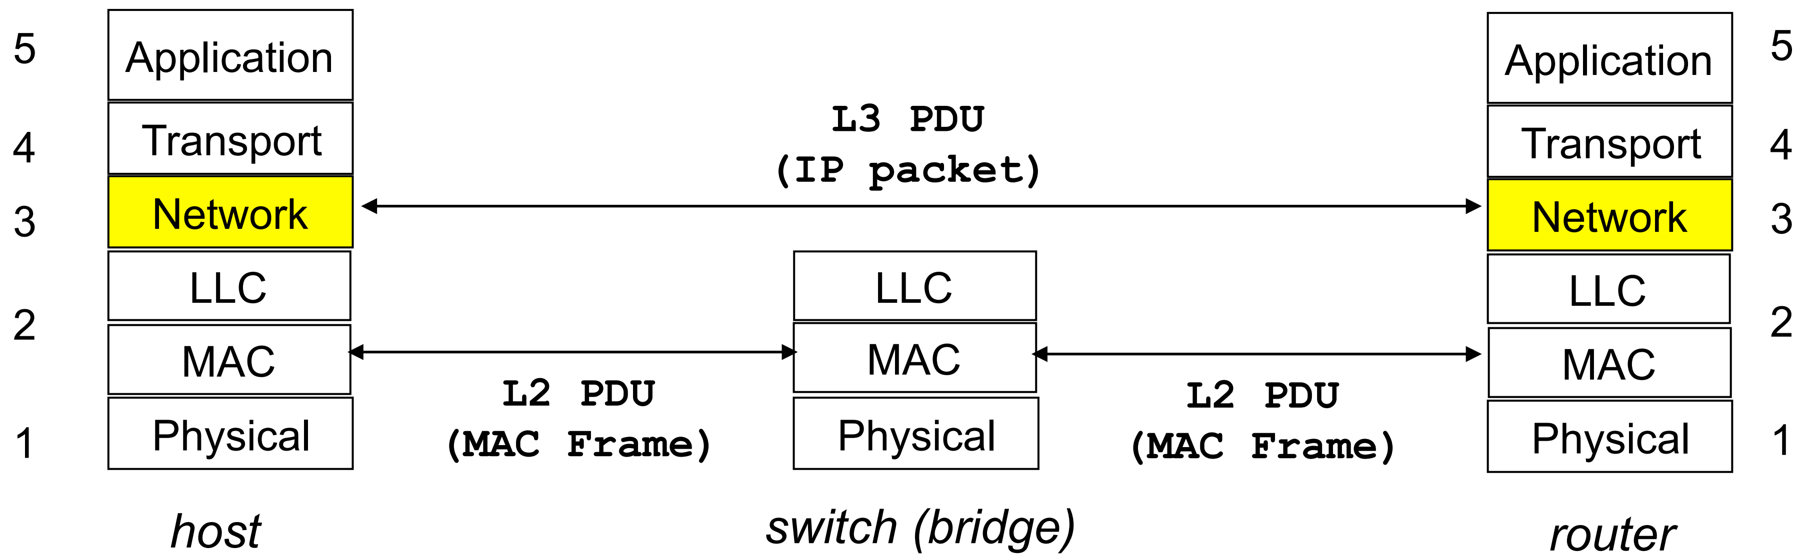

- Routers are layer 3 intermediate systems
- Explicit forwarding
  - host has to know the address of the first router
- Management protocols (control, routing, configuration)

Protocols

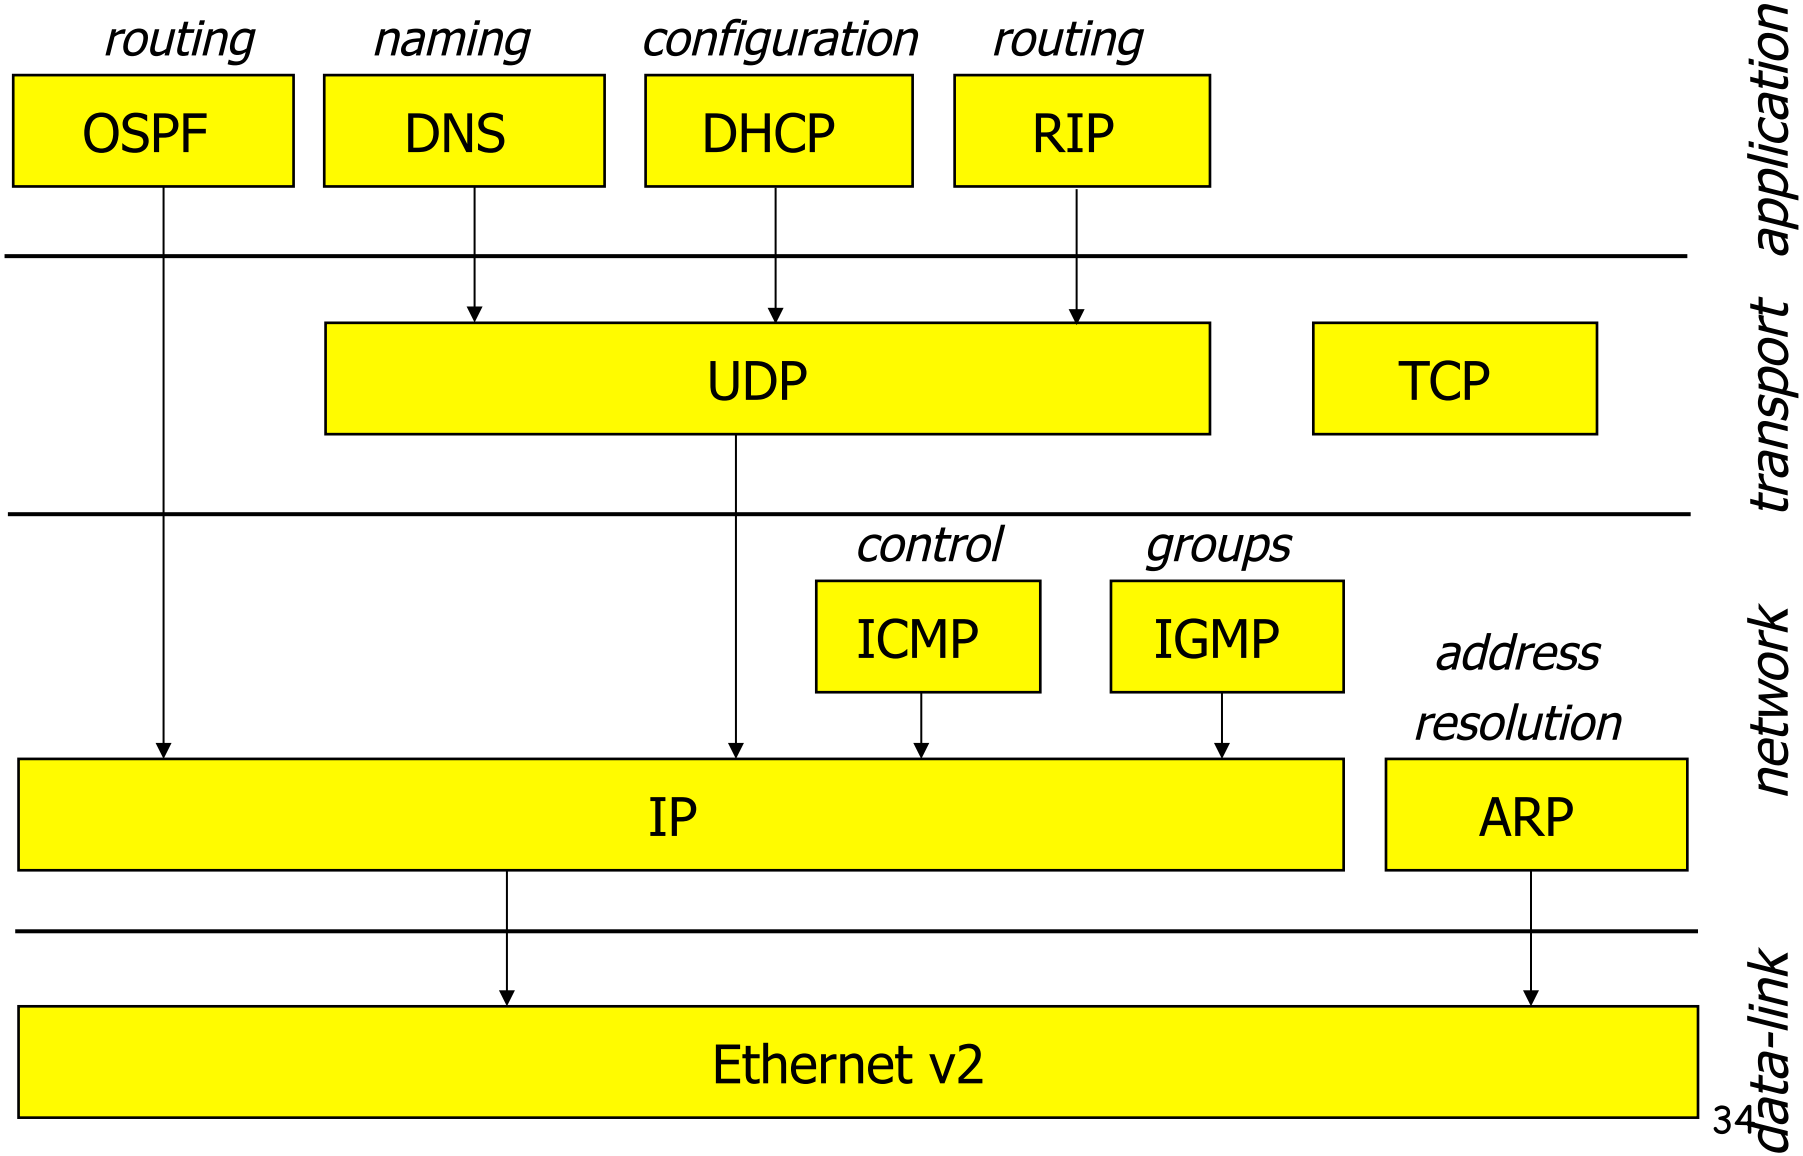
routing
OSPF
naming
DNS
configuration
DHCP
routing
RIP
application
UDP
TCP
transport
control
ICMP
groups
IGMP
address
resolution
network
IP
ARP
Ethernet v2
data-link

Autonomous systems

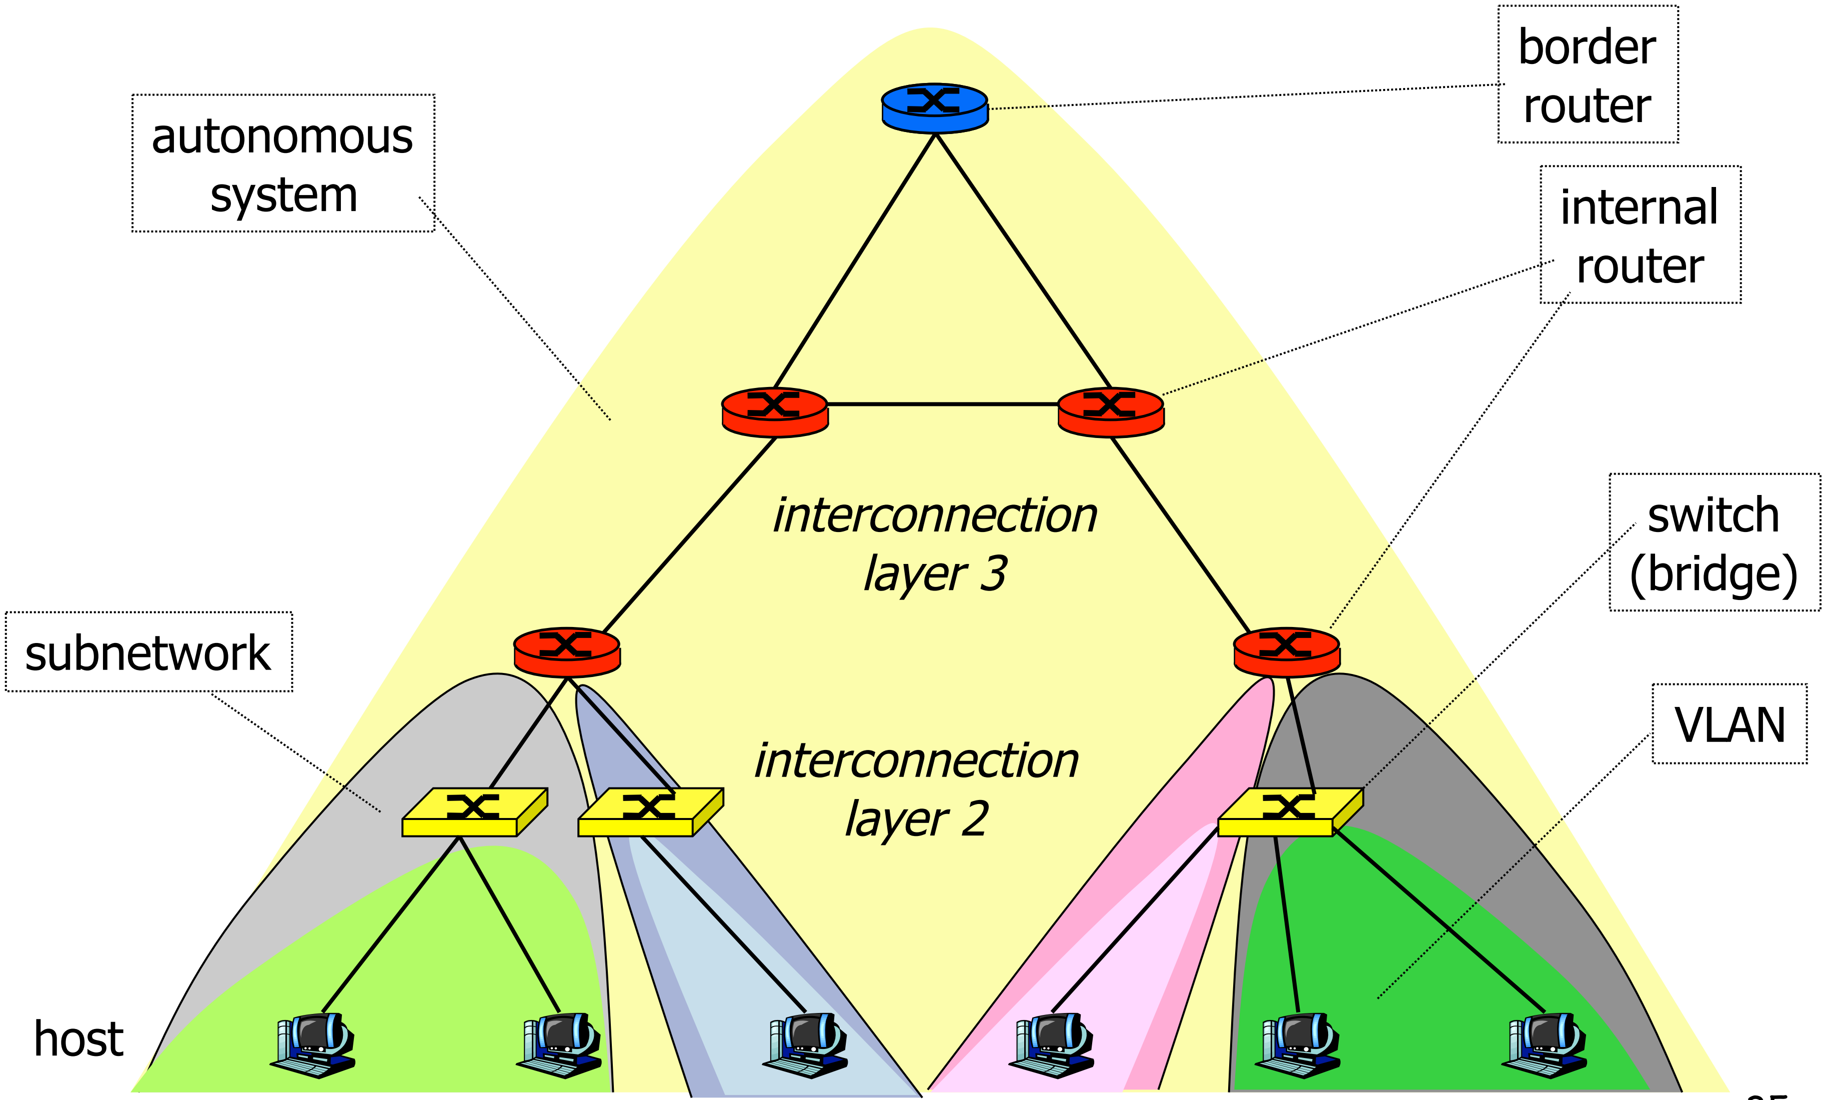

Internet

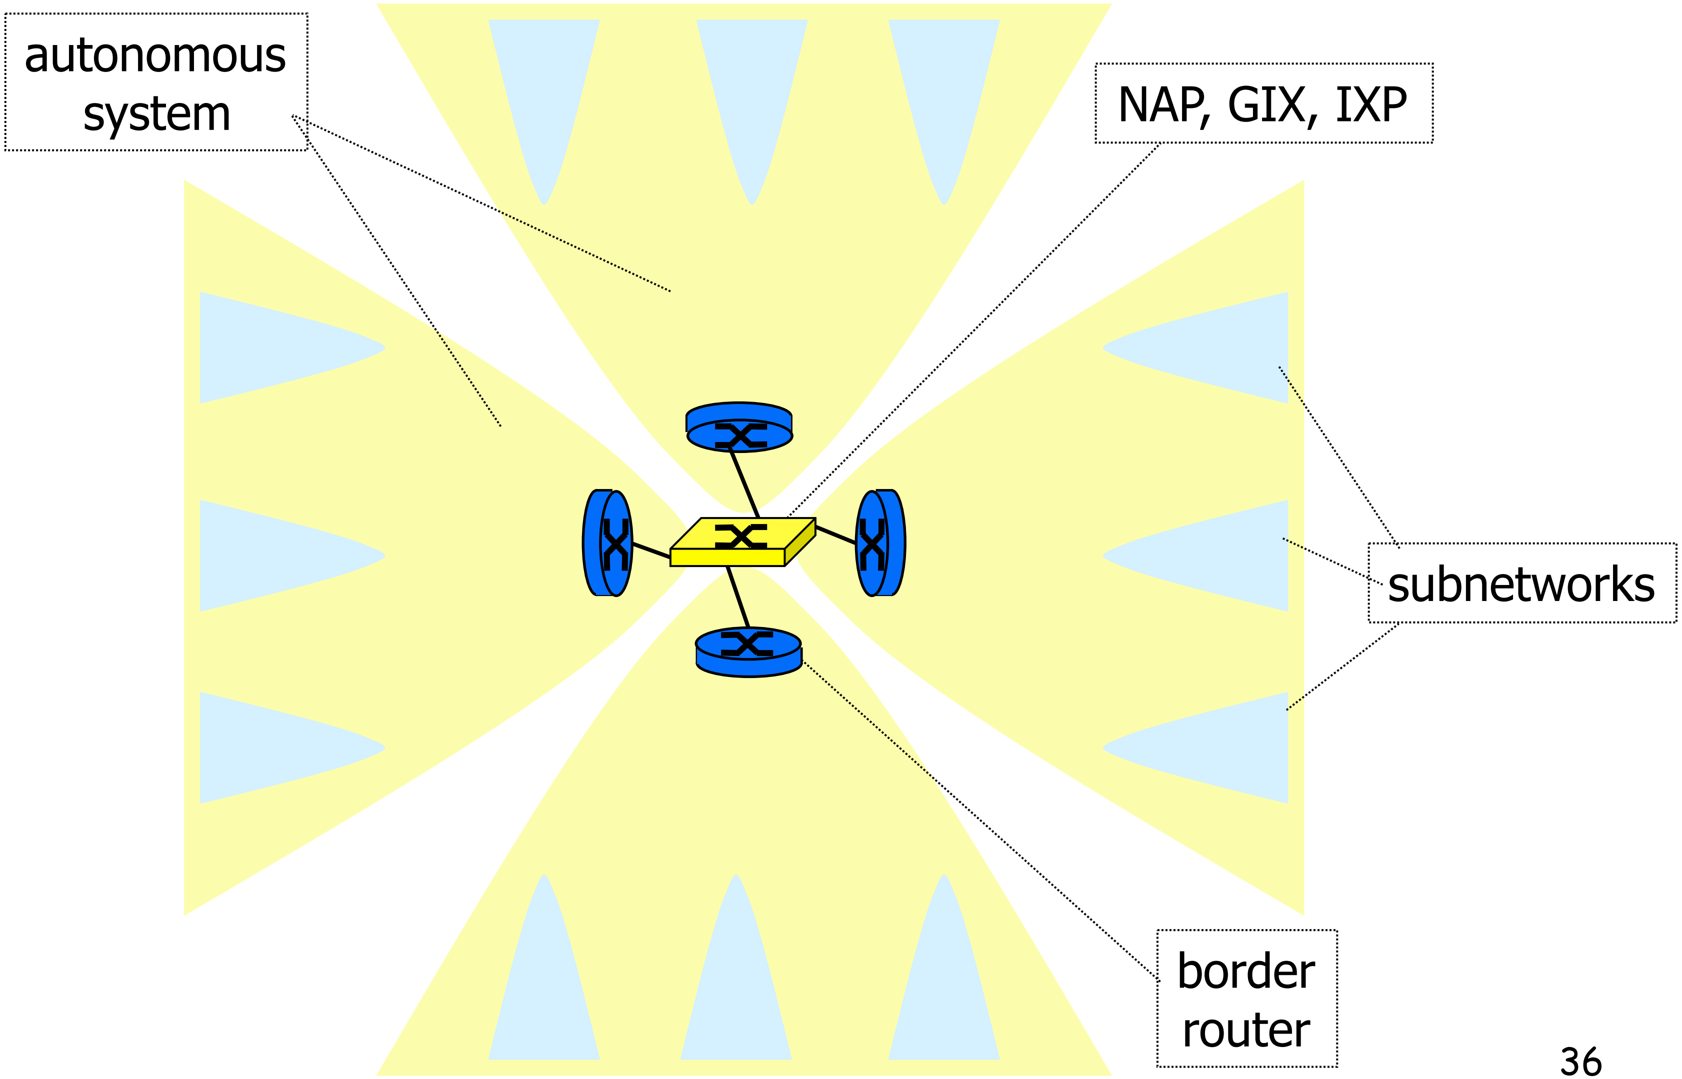

Interconnection of AS

- Border routers
  - interconnect AS
- NAP or GIX, or IXP
  - exchange of traffic - peering
- Route construction
  - based on the path through a series of AS
  - based on administrative policies
  - routing tables: aggregation of entries
  - works if no loops and at least one route - routing protocols (EGP - External Routing Protocols)

Protocols

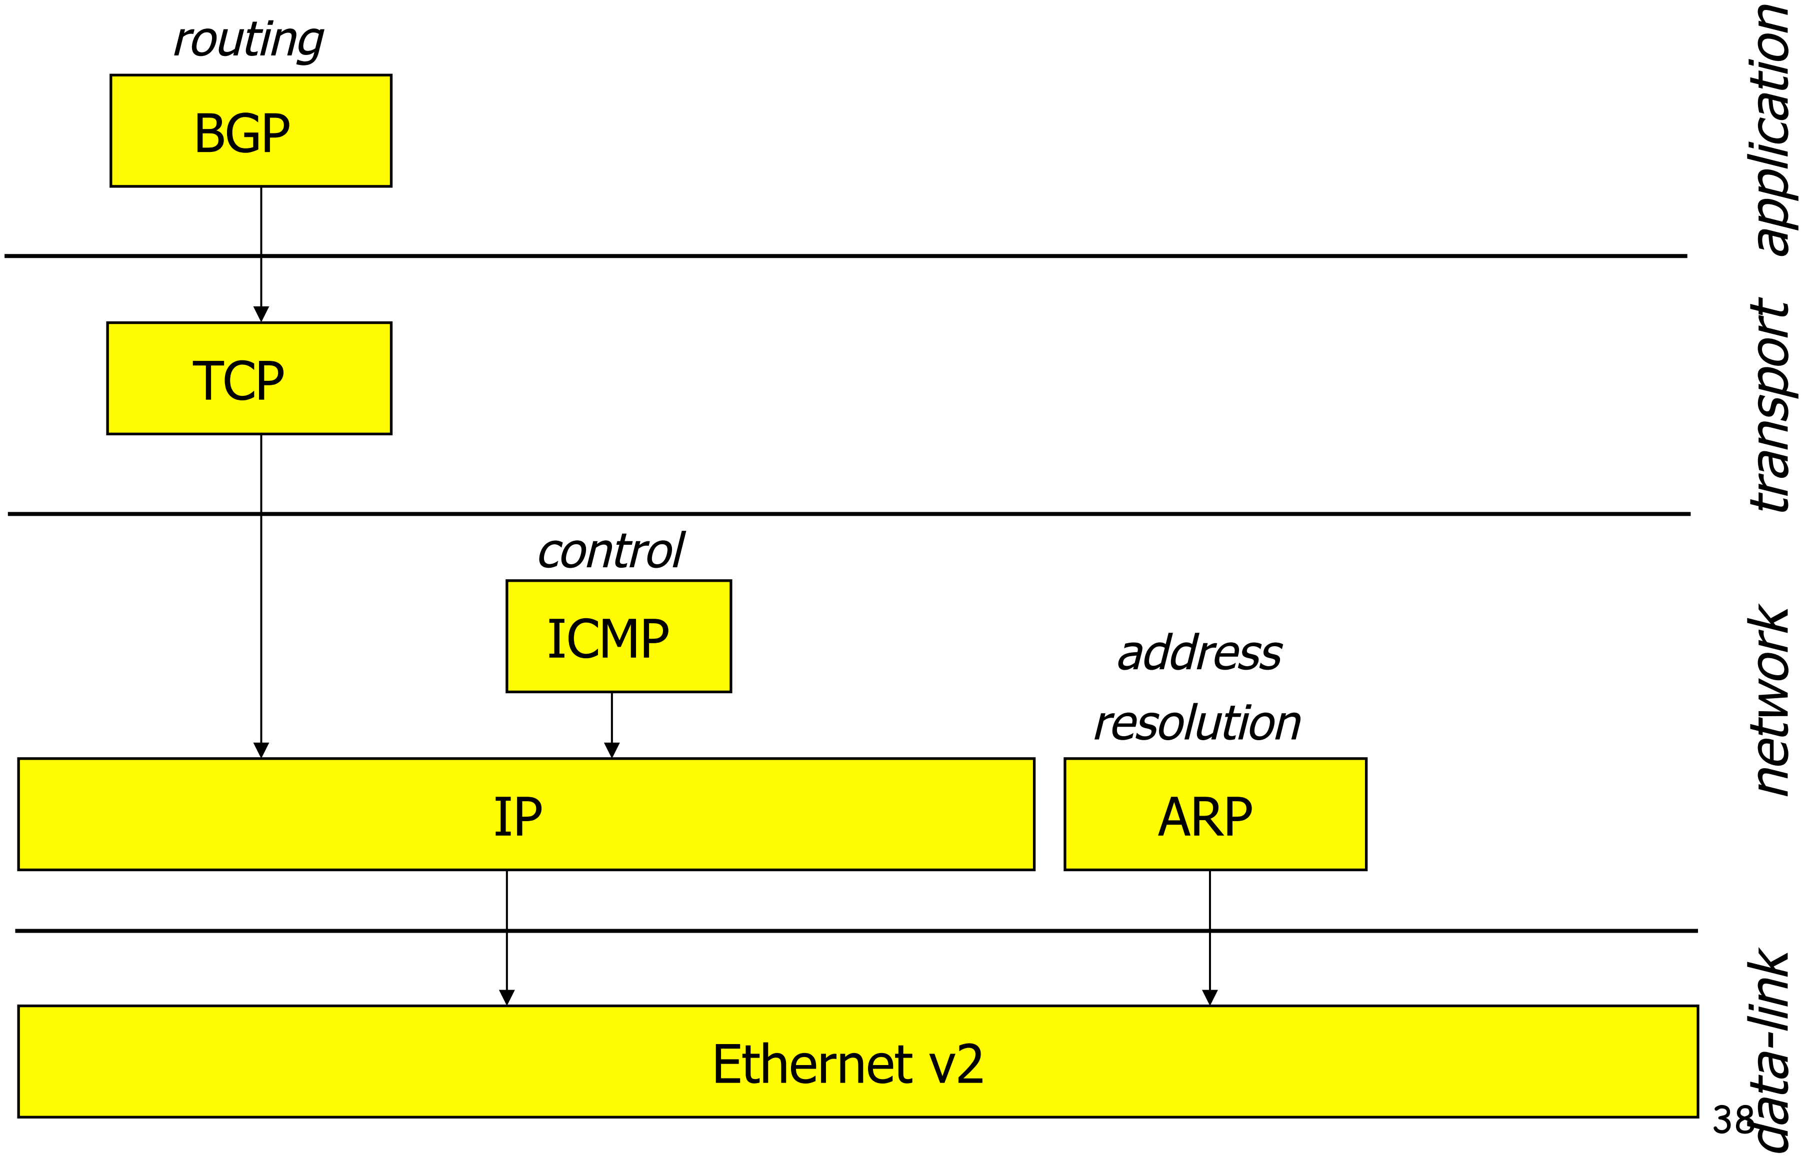
routing
BGP
TCP
control
ICMP
address
resolution
IP
ARP
Ethernet v2
application
transport
network
data-link

Performance

- Bit Rate (débit binaire) of a transmission system
  - bandwidth, throughput
  - number of bits transmitted per time unit
  - units: b/s or bps, kb/s = 1000 b/s, Mb/s = 10e+06 b/s, Gb/s=10e+09 b/s

- Latency or Delay
  - time interval between the beginning of a transmission and the end of the reception
  - RTT - Round-Trip Time

Delay in packet-switched networks

packets experience delay on end-to-end path

- four sources of delay at each hop

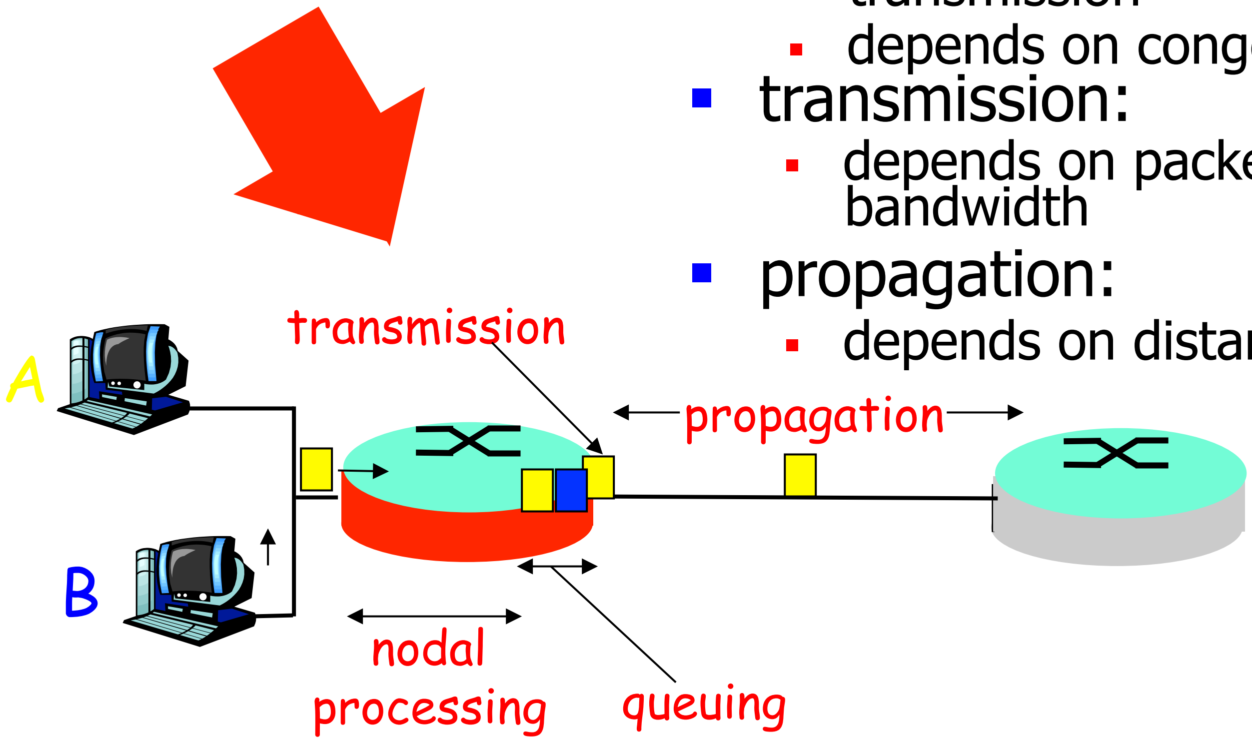

- nodal processing:
  - check bit errors
  - determine output link
- queuing
  - time waiting at output link for transmission
  - depends on congestion level of node
- transmission:
  - depends on packet length and link bandwidth
- propagation:
  - depends on distance between nodes

Delay

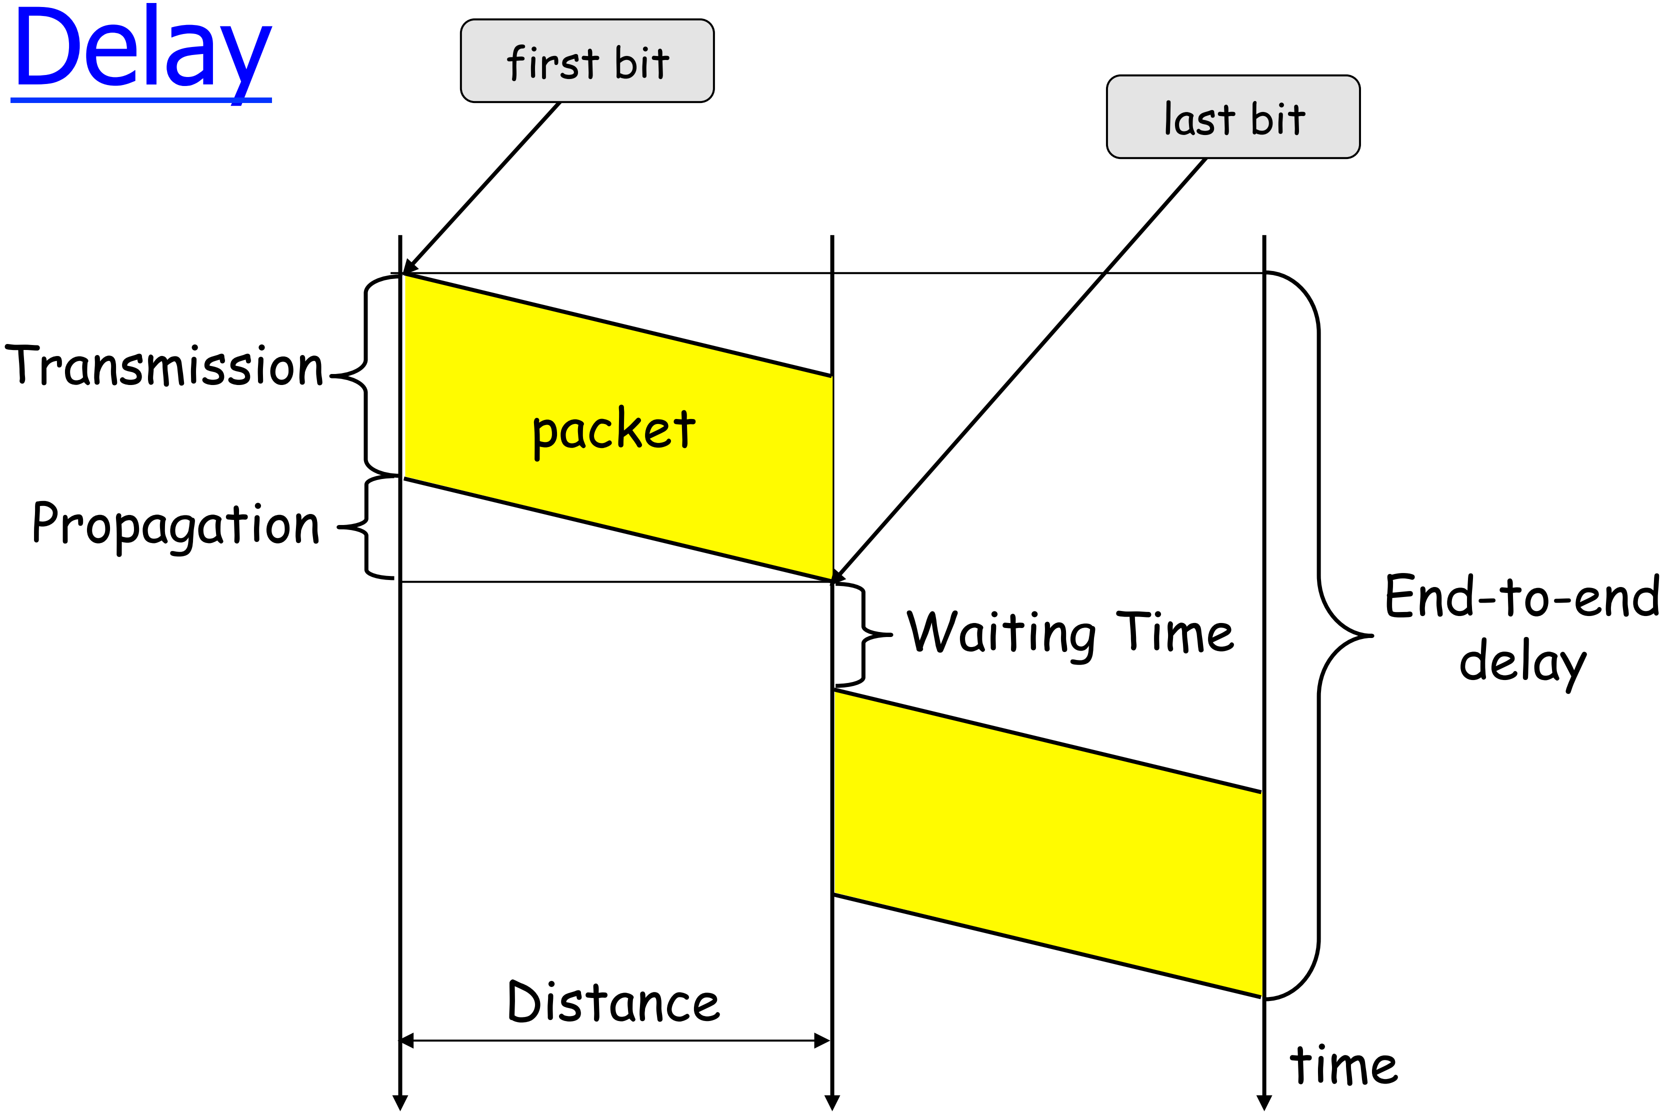
first bit
last bit
Transmission
packet
Propagation
Waiting Time
End-to-end
delay
Distance
time

Performance

- Latency
  - Latency = Propagation + Transmission + Wait
  - Propagation = Distance / Speed
    - copper : Speed = 2.3×10^8 m/s
    - glass : Speed = 2×10^8 m/s
    - Transmission = Size / BitRate
- 5 μs/km
- New York - Los Angeles in 24 ms
  - request - 1 byte, response - 1 byte: 48 ms
  - 25 MB file on 10 Mb/s: 20 s
- World tour in 0.2 s

Example

- At time 0, computer A sends a packet of size 1000 bytes to B; at what time is the packet received by B (speed = 2e+08 m/s)?

distance	20 km	20000 km	2 km	20 m
bit rate	10kb/s	1 Mb/s	10 Mb/s	1 Gb/s
propagation	?	?	?	?
transmission	?	?	?	?
latency	?	?	?	?
	modem	*satellite*	*LAN*	*Hippi*

Example

- At time 0, computer A sends a packet of size 1000 bytes to B; at what time is the packet received by B (speed = 2e+08 m/s)?

distance	20 km	20000 km	2 km	20 m
bit rate	10kb/s	1 Mb/s	10 Mb/s	1 Gb/s
propagation	0.1ms	100 ms	0.01 ms	0.1μs
transmission	800 ms	8 ms	0.8 ms	8 μs
latency	800.1 ms	108 ms	0.81 ms	8.1 μs
	modem	*satellite*	*LAN*	*Hippi*

Waiting time

- Queuing system M/M/1
  - Inter-arrival times ~ exponentially distributed
  - service times ~ exponentially distributed
  - arrival rate λ, service rate μ, utilization $\rho = \lambda/\mu$
  - number of packets N, waiting time T

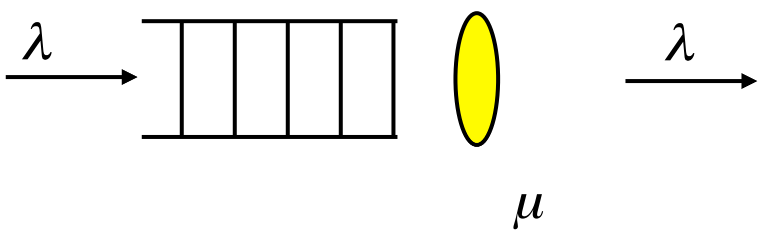

$$N = \frac{\rho}{(1-\rho)}$$

$$T = \frac{1}{\mu(1-\rho)}$$

$$T = \frac{N}{\lambda}$$

Waiting time

- Average packet length 1500 bytes
  - link with 1 Mb/s bit rate (propagation = 0)
    - transmission time 12 ms
    - service rate 83 packet/s

λ	[p/s]	10	40	60	70	
$1/\lambda$	[ms]	100	25	16	14	
T	[ms]		13	23	43	76

Waiting time

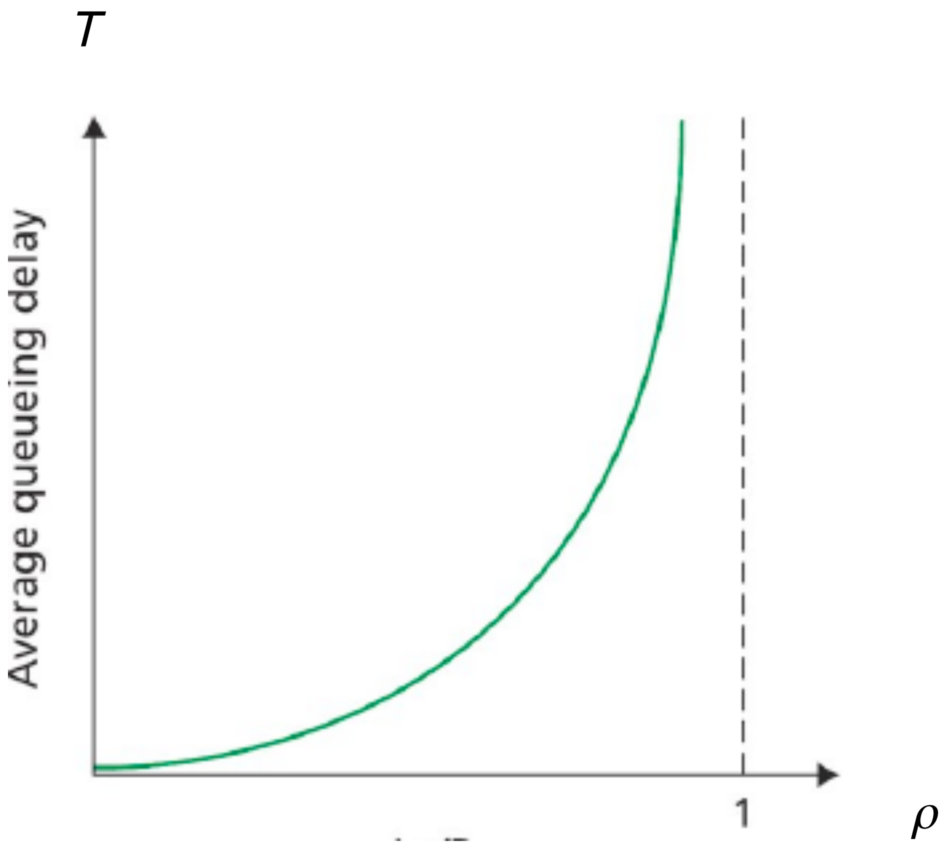

Summary

- Network architectures
  - protocol architectures
    - different protocol stacks, overlaid stacks
  - interconnection structure
    - switches, routers
  - related protocols
    - complex protocol families
- Performance
  - transmission
  - propagation
  - queuing delay

Problem

- Two hosts are connected via the Internet through 9 routers. The distance between the hosts is 10000 km. All the links in the network have 100 Mb/s bit rate.
- The user working on host A downloads a page of 1000 bytes from server B. What is the total time of the download (between the click and the instant when the page is downloaded)?
- Assume that:
  - HTTP uses a TCP connection with the MTU of 1460 bytes. Other TCP parameters are supposed to be known.
  - We ignore processing and waiting times, as well as the transmission time of short segments (short means that they are less than 1000 bytes), for instance connection establishment segments, ACKs, HTTP request. We also ignore the HTTP header attached to the page contents.

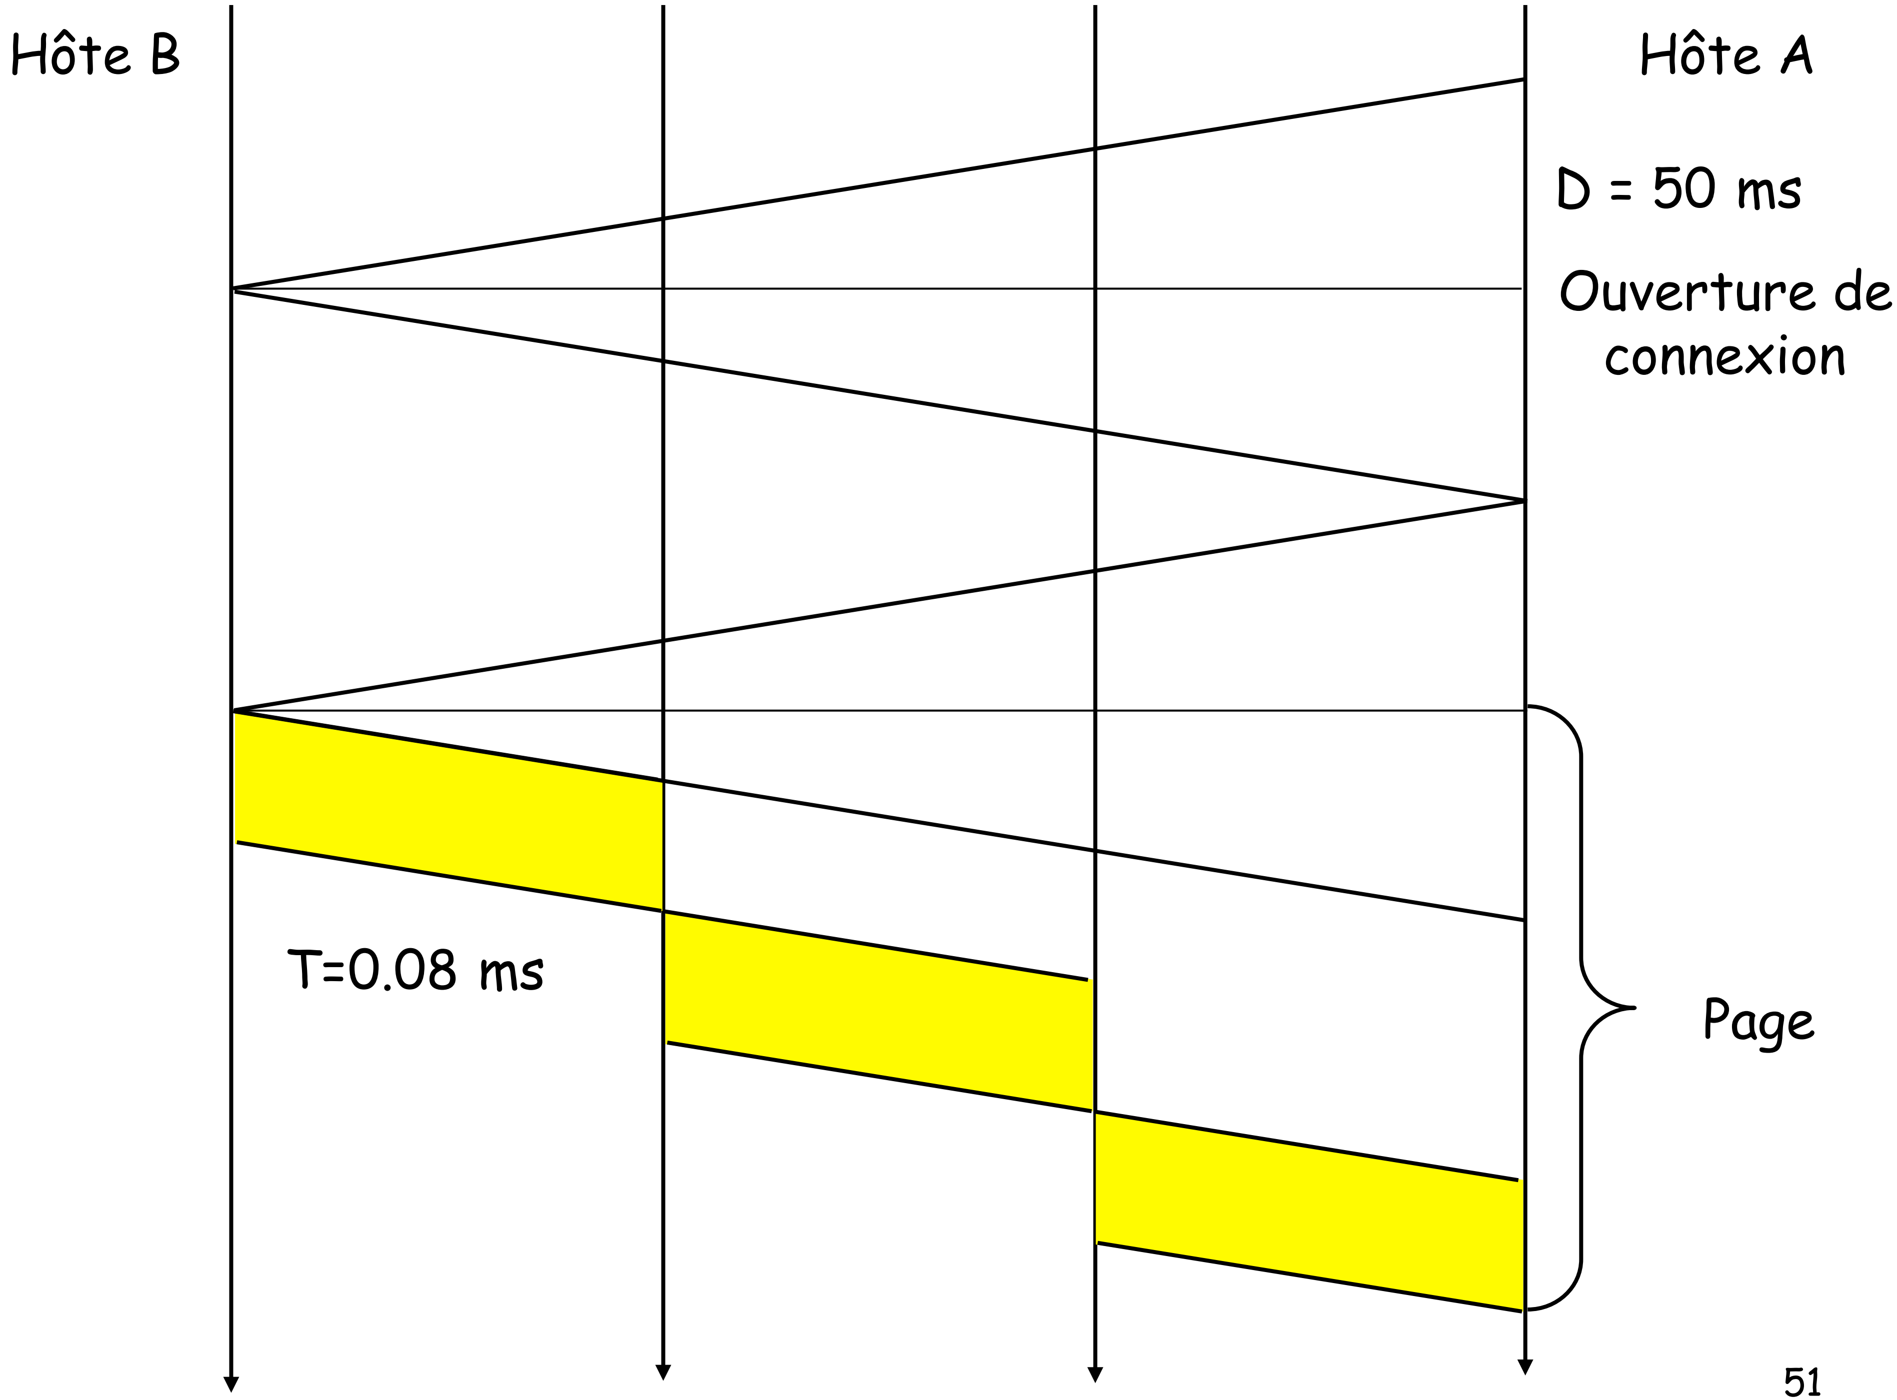
Hôte B
Hôte A
D = 50 ms
Ouverture de connexion
T=0.08 ms
Page

Solution

- Propagation = distance / speed = 50 ms
- Transmission = size / bitrate = 0.08 ms

- délai = 3 x 50 ms + 50 ms + 10 x 0.08 = 200.8 ms

- débit : 1000 x 8 bits / 200.8 ms = 39 kb/s