Performance

Performance - Motivating example

- Consider this real-life example of a large bank with headquarters in **Europe** and operations in **North America**.
- Problem: a business unit with European users trying to access an important application from across the pond.
- **Performance was horrible (response time).**
- CIO ordered his trusted network operations manager to **fix the problem**. The network manager dutifully investigated, measuring the transatlantic link utilization and router queue statistics: no problems with the network, as it was only **3 percent utilized**.
- "I don't care, **double the bandwidth**!" the CIO ordered. The network manager complied, installing a **second OC-3 link**. And, guess what?
- The network went from **3 percent to 1.5 percent utilized**, and **application performance was still horrible**. That CIO didn't *know jack about network performance*.

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
  - OC3/STM1 - 155 Mb/s, OC12/STM4 - 622 Mb/s, and OC48/STM-16 - 2.5 Gb/s, OC192/STM-48 10 Gb/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

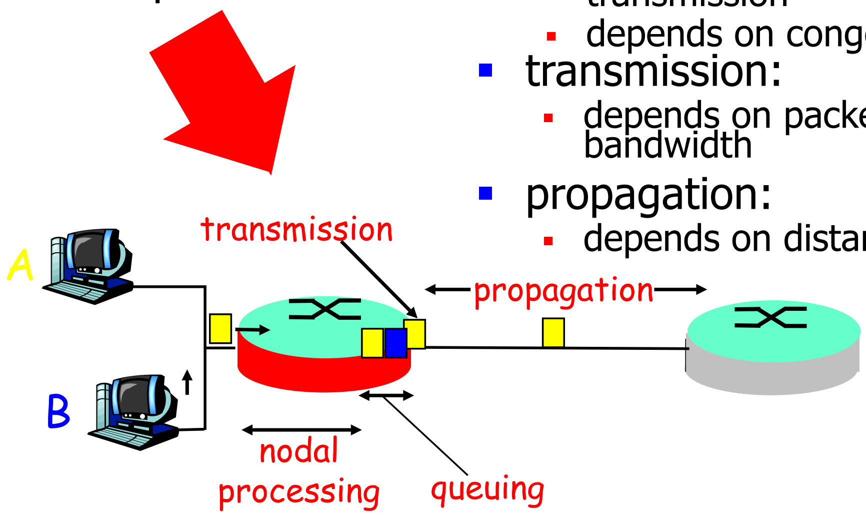

Delay

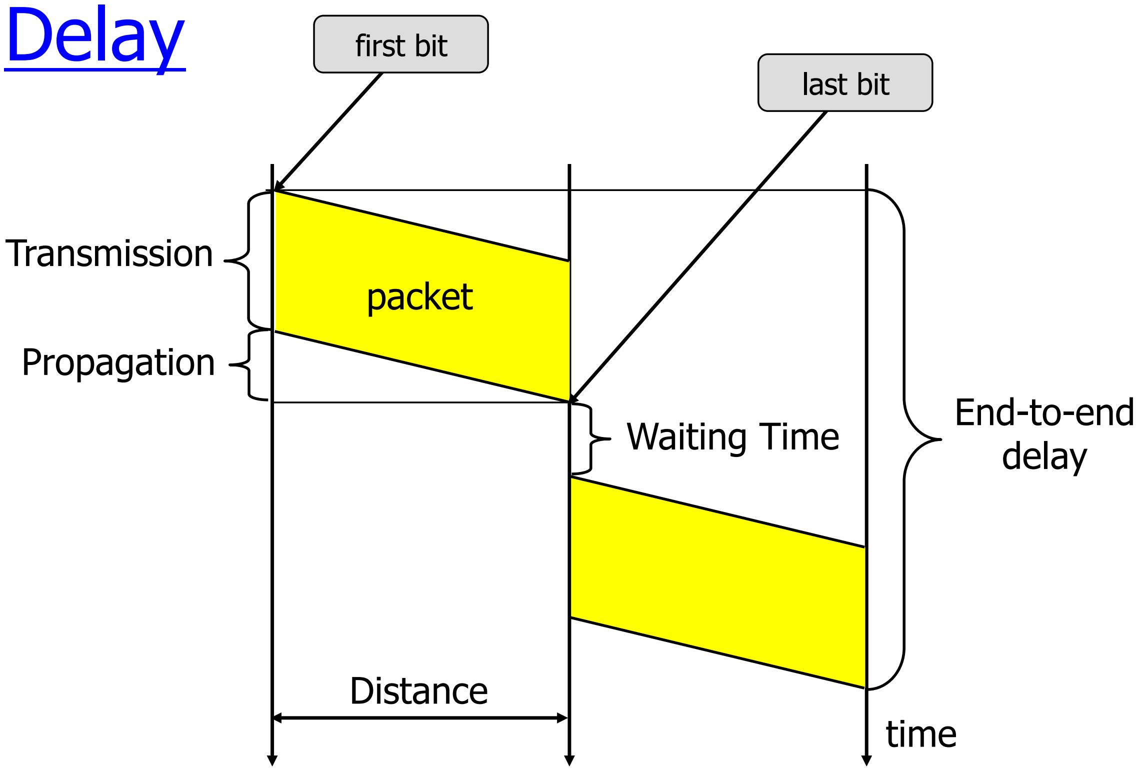

Performance

- 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
- Around the world 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	0.1ms	100 ms	0.01 ms	0.1μs
transmission	800 ms	8 ms	0.8 ms	8 μs
latency	?	?	?	?
	modem	*satellite*	*Cable*	*LAN*

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	[redacted]	[redacted]	[redacted]	[redacted]
	modem	*satellite*	*Cable*	*LAN*

Bandwidth-Delay Product

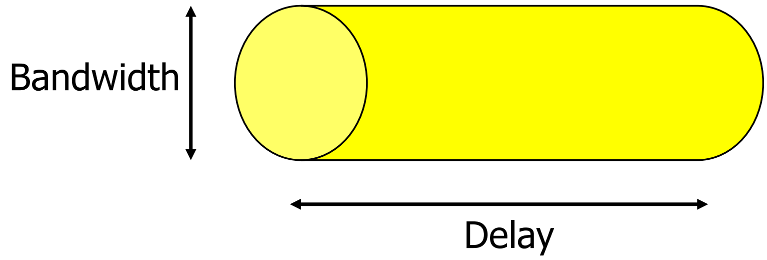

- Bandwidth-Delay product
  - how many bits should we send before the arrival of the first bit?
  - good utilization - keep the pipe filled!

A Simple Protocol: Stop and Go

- Packets may be lost during transmission:
  bit errors due to channel imperfections, various noises.
- Computer A sends packets to B; B returns an acknowledgement packet immediately to confirm that B has received the packet;

  A waits for acknowledgement before sending a new packet; if no acknowledgement comes after a delay *T1*, then A retransmits

A Simple Protocol: Stop and Go

- **Question**: What is the maximum throughput assuming that there are no losses?

  notation:
  - packet length = L, constant (in bits);
  - acknowledgement length = l, constant
  - channel bit rate = b;
  - propagation = D
  - processing time = 0

Solution (1)

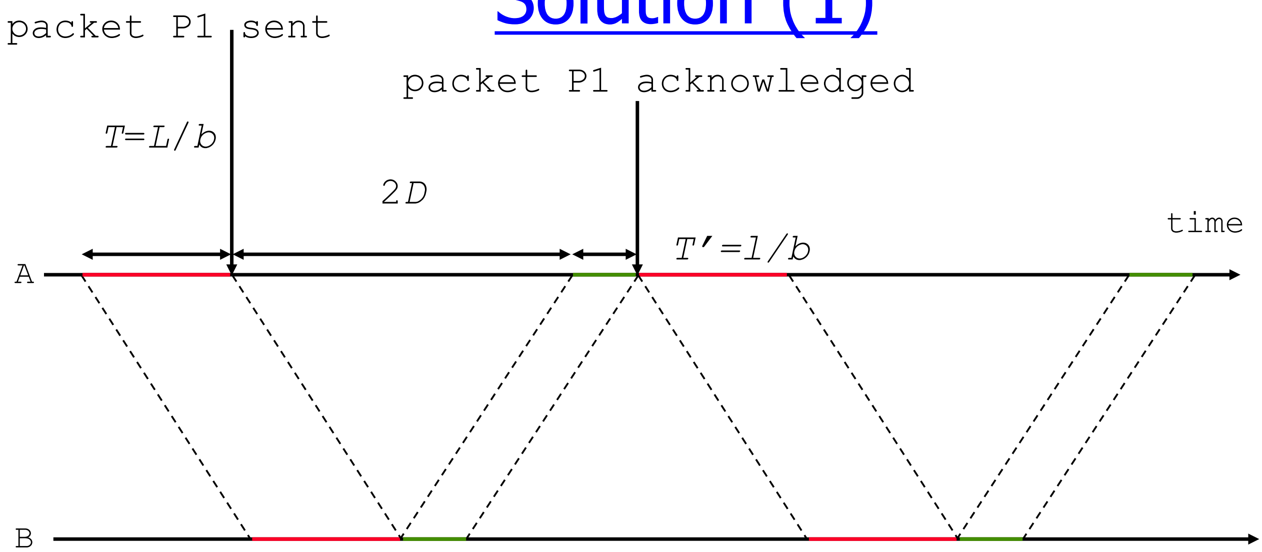

cycle time = $T + 2D + T'$

useful bits per cycle time = L

throughput = $L / (L/b + l/b + 2D) = b / (\omega + \beta/L)$

with $\omega = (L+l)/L$=overhead and $\beta = 2Db$=**bandwidth-delay product**

Bandwidth delay

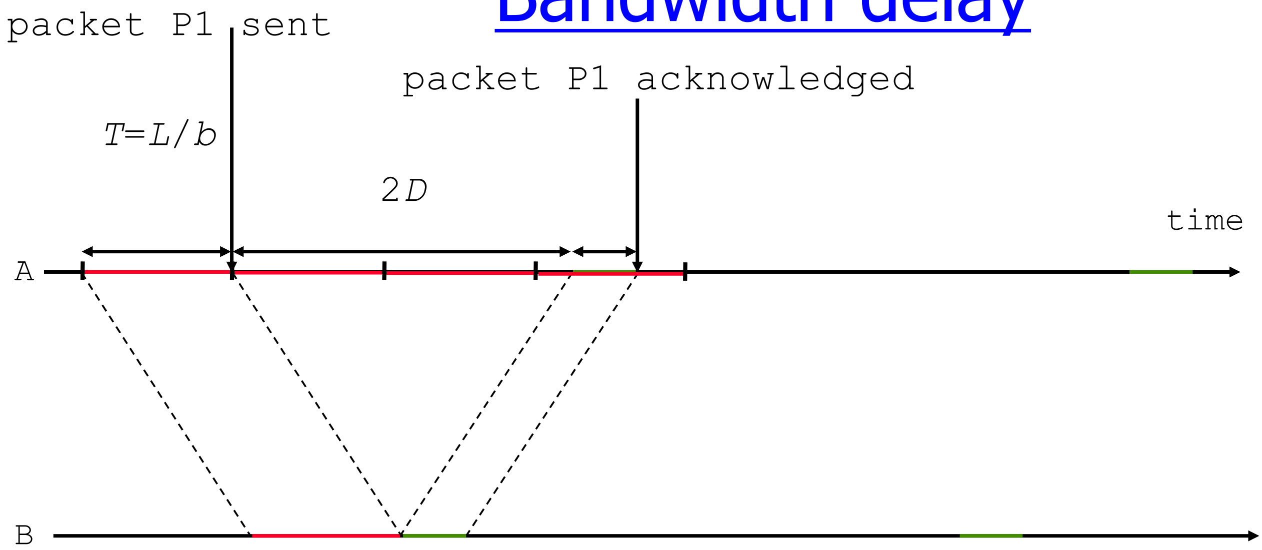

window in time = $T + 2D + T'$

window in bits = $(T + 2D + T')b = L+l + \beta$

Solution (2)

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
reception time	800.1 ms	108 ms	0.81 ms	8.1 μs
	modem	*satellite*	*Cable*	*LAN*
$\beta = 2Db$	2 bits	200 000 bits	200 bits	200 bits
throughput = *b x*	99.98%	3.8%	97.56%	97.56%

Summary

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

Waiting time

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

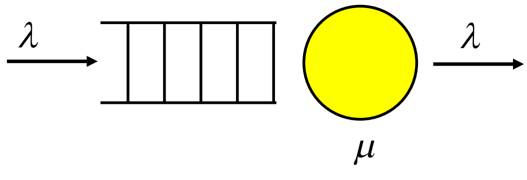

$$N = \frac{\rho}{(1-\rho)} \qquad T = \frac{1}{\mu(1-\rho)} \qquad T = \frac{N}{\lambda}$$

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

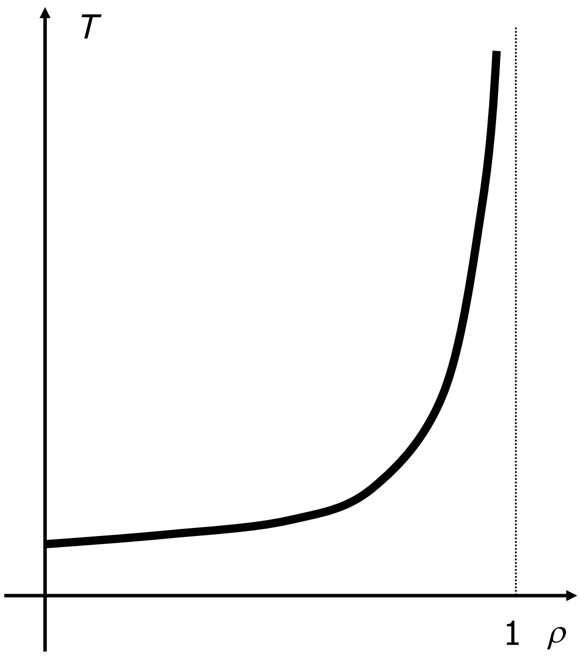

Exercice

- Two hosts are connected via the Internet through 19 routers. The distance between the hosts is 10000 km. All the links in the network have 10 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.

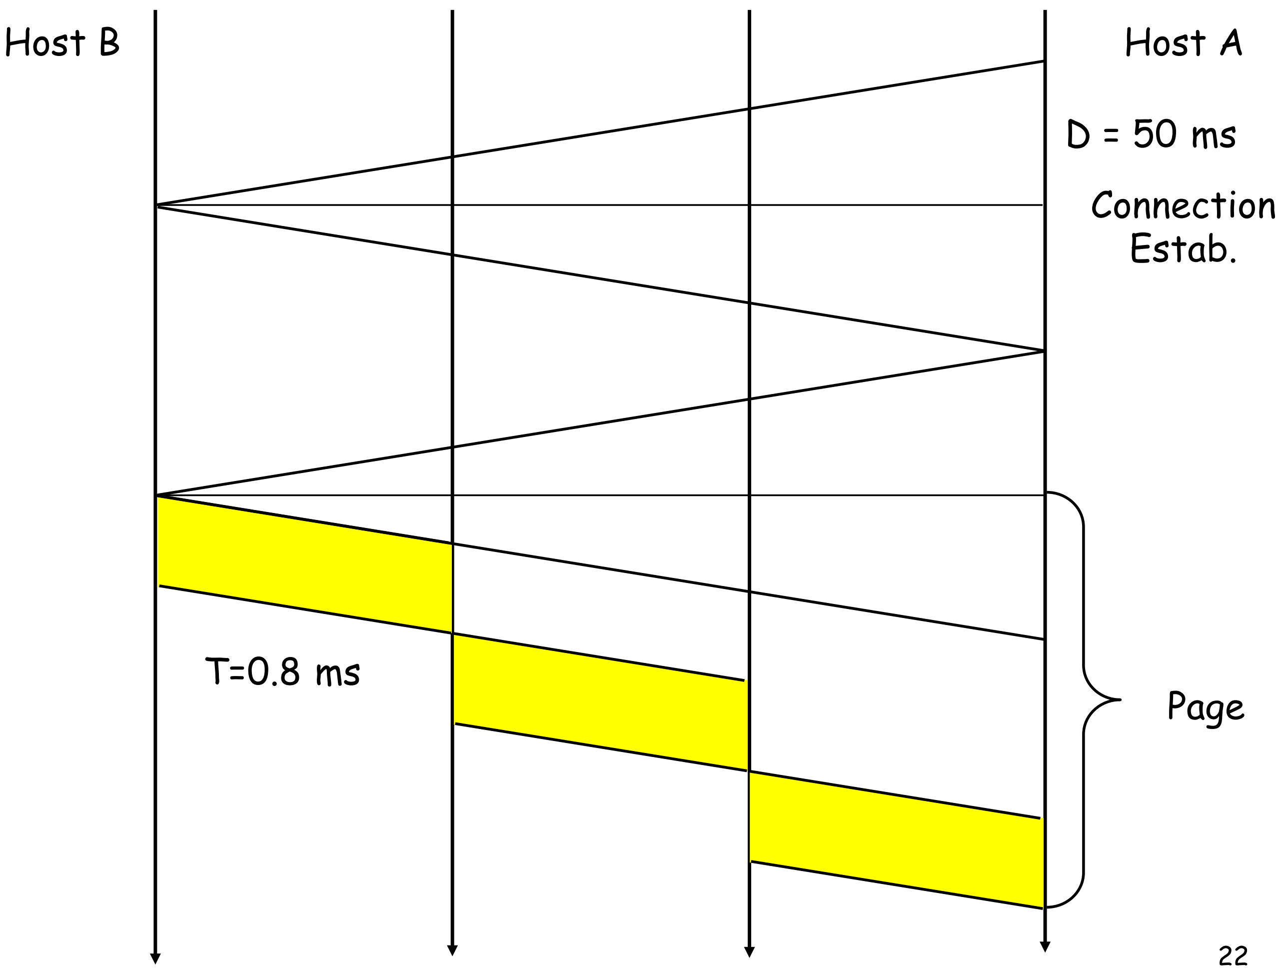

Residential access

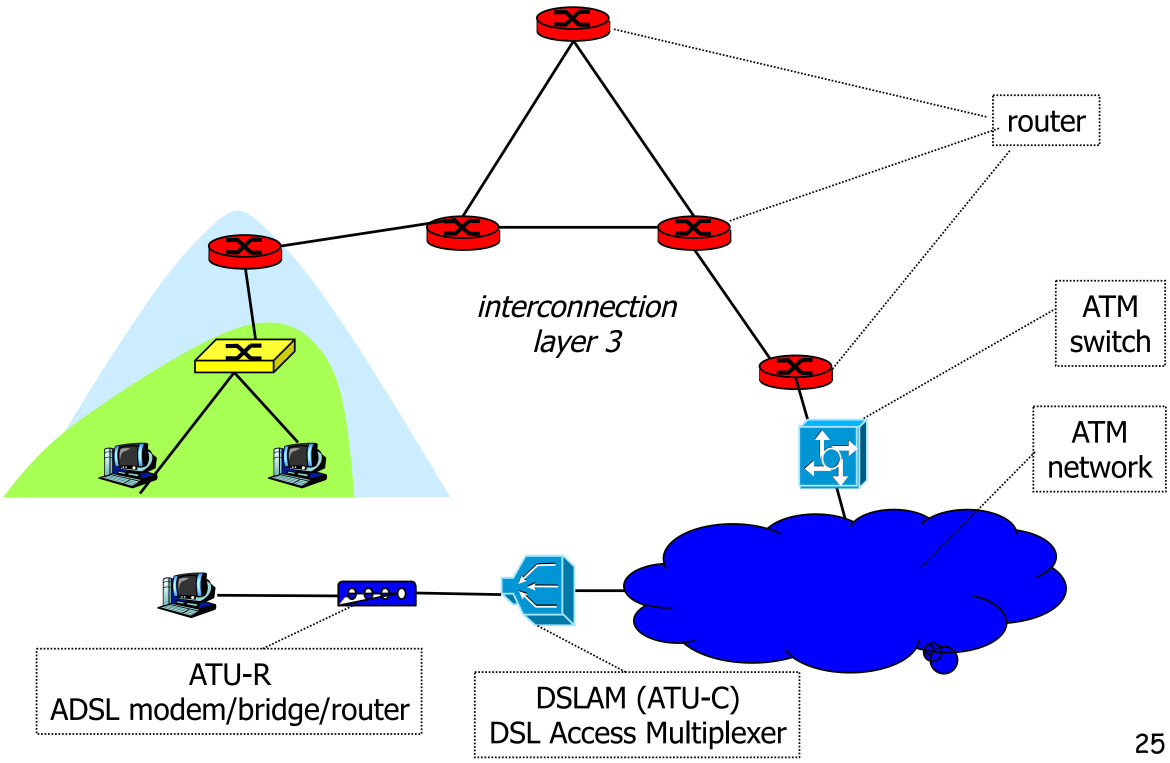

Protocol architecture

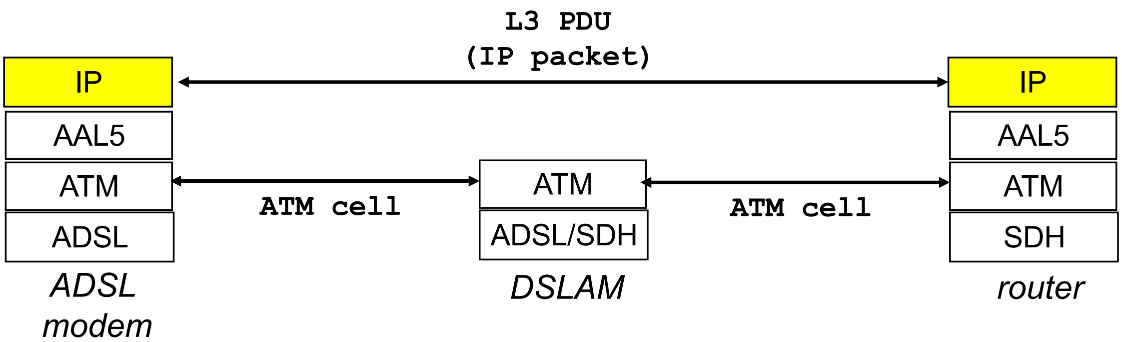

- IP over ATM
  - requires fixed IP address

Protocol architecture

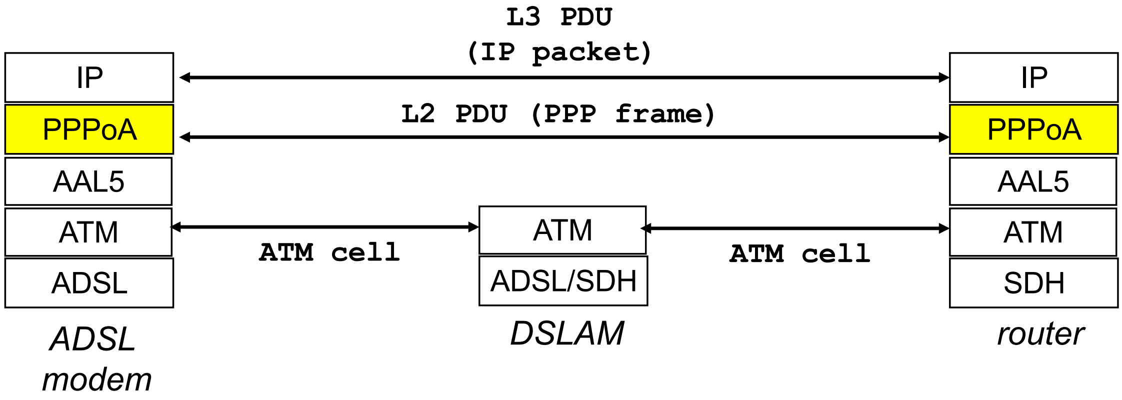

- PPP over ATM (PPPoA)
  - multiple users share ADSL link

Protocol architecture

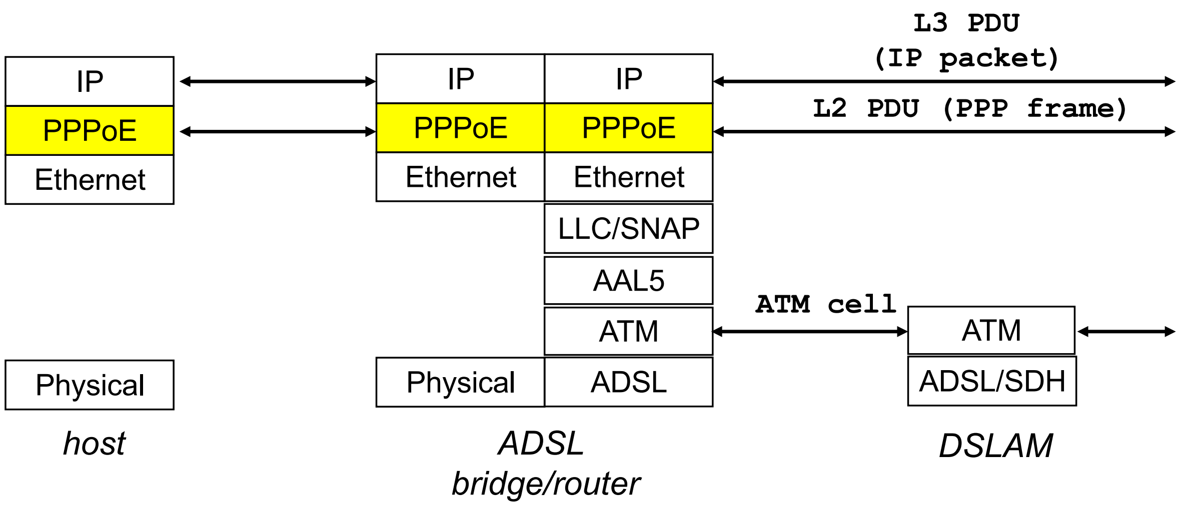

- PPP over Ethernet (PPPoE)
  - multiple users share ADSL link