

CSCE 463/612

Networks and Distributed Processing

Spring 2018

Introduction II

Dmitri Loguinov

Texas A&M University

January 25, 2018

Original slides copyright © 1996-2004 J.F Kurose and K.W. Ross

Chapter 1: Roadmap

1.1 What *is* the Internet?

1.2 Network edge

1.3 Network core

1.4 Network access and physical media

1.5 Internet structure and ISPs


1.6 Delay & loss in packet-switched networks

1.7 Protocol layers, service models

1.8 History

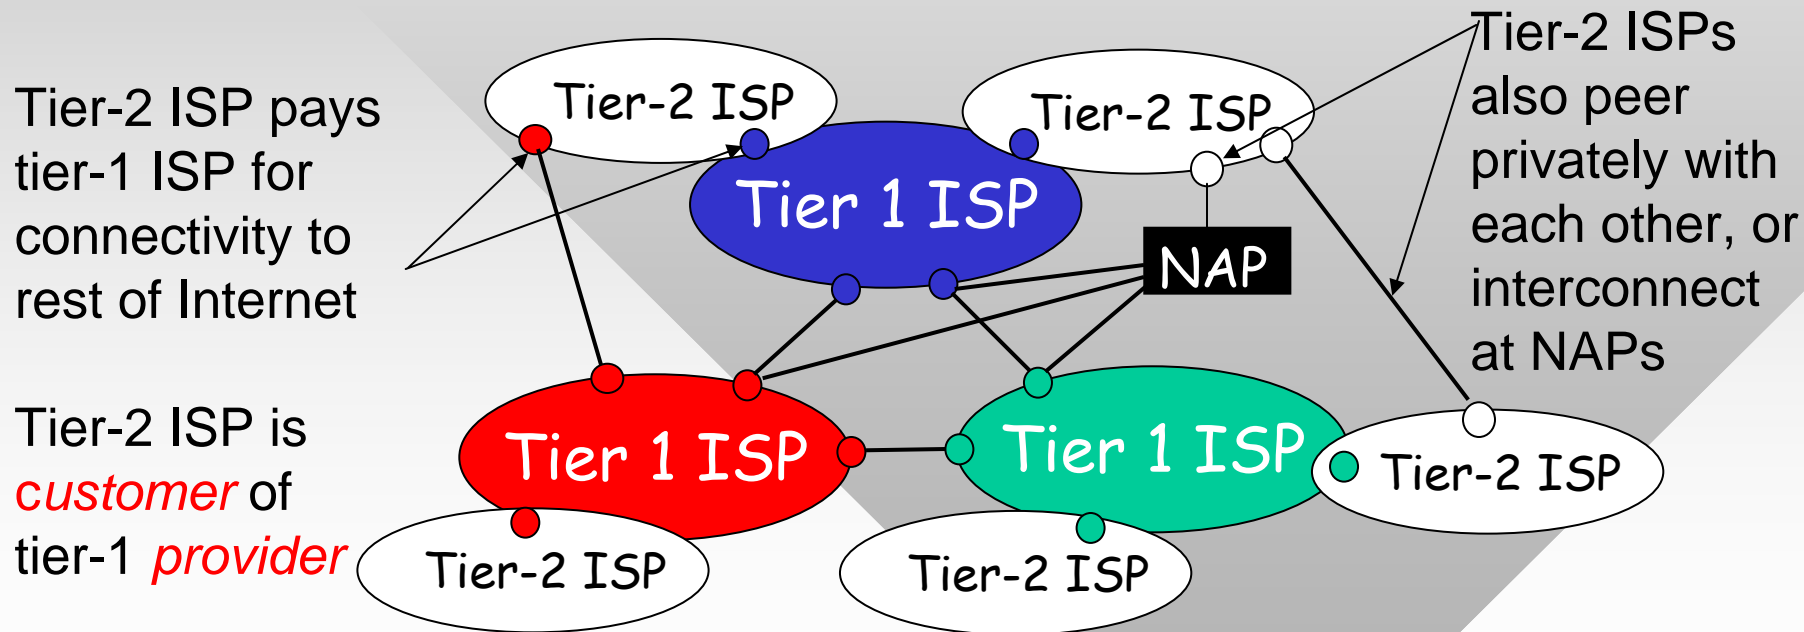
Internet: Network of Networks

- Roughly hierarchical
 - **In the center:** “tier-1” ISPs (e.g., Level3, Sprint, AT&T, Verizon), national/international coverage
 - Treat each other as equals, do not pay for upstream bandwidth
 - Form the **backbone** of the Internet



Internet: Network of Networks

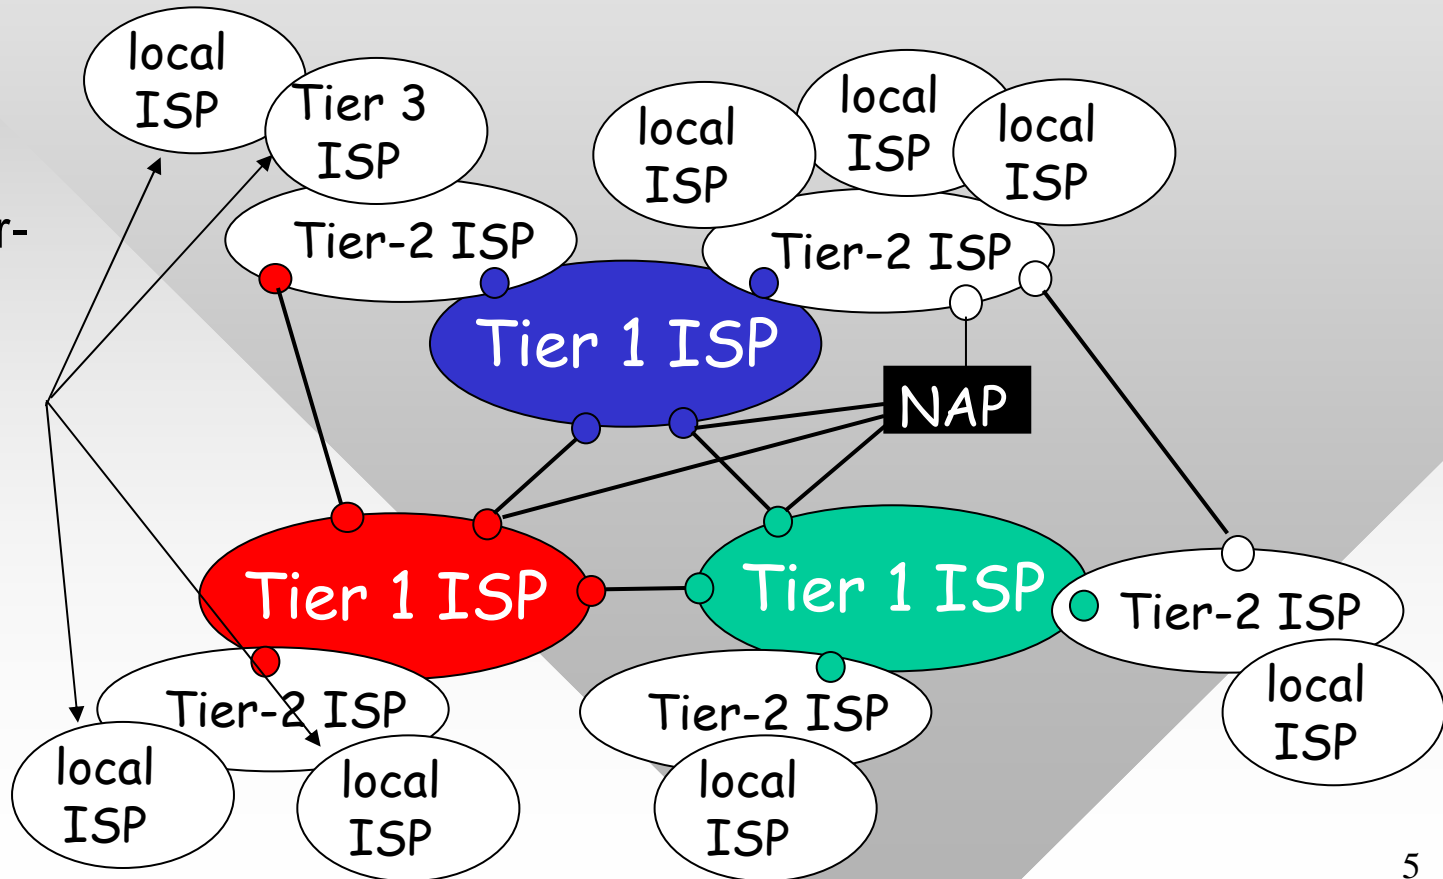
- “Tier-2” ISPs: smaller (often regional) ISPs
 - Connect to one or more tier-1 ISPs, possibly other tier-2 ISPs



Internet Structure: Network of Networks


- “Tier-3” ISPs and local ISPs
 - Last hop (“access”) network (closest to end systems)

Local and tier-3 ISPs are customers of higher tier ISPs connecting them to rest of Internet



Internet Structure: Network of Networks

- A packet passes through many networks!



Chapter 1: Roadmap

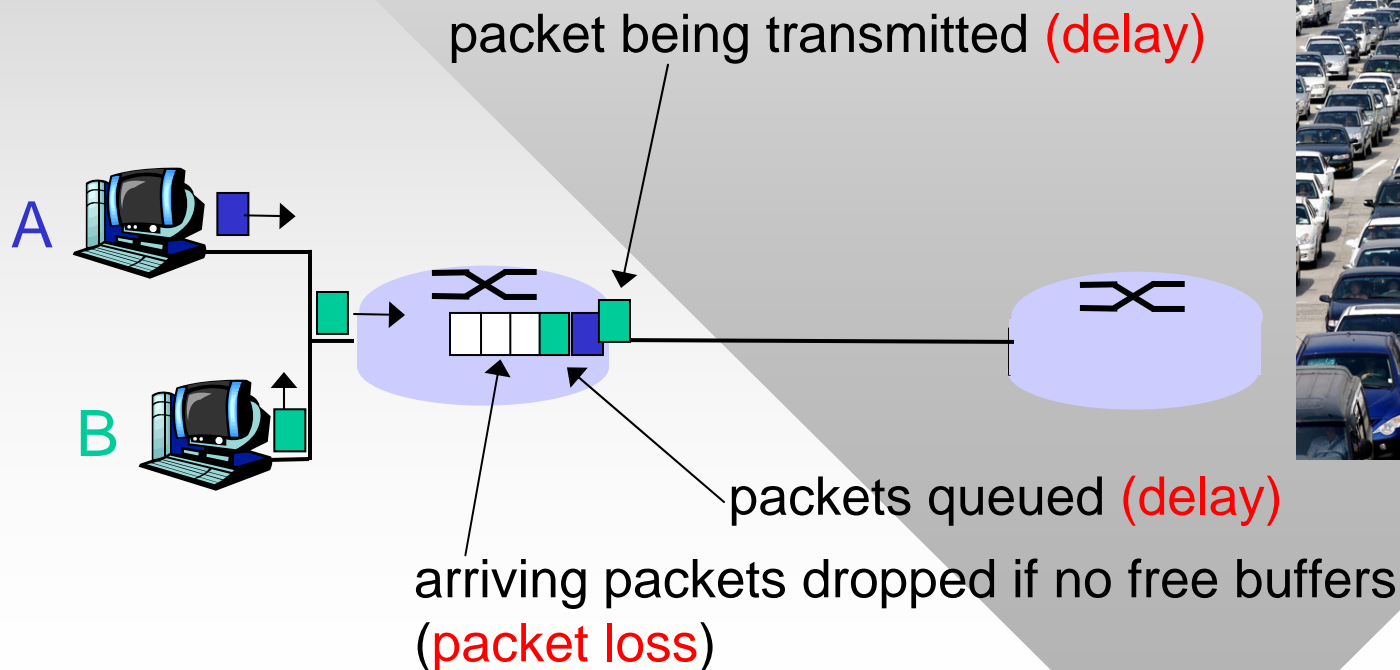
- 1.1 What *is* the Internet?
- 1.2 Network edge
- 1.3 Network core
- 1.4 Network access and physical media
- 1.5 Internet structure and ISPs
- 1.6 Delay & loss in packet-switched networks
- 1.7 Protocol layers, service models
- 1.8 History

How Do Loss and Delay Occur?

Packets *queue* in router buffers (typically FIFO queues)

- **If packet arrival rate exceeds output link capacity:**

- Packets queue, wait for their turn
- Analogy: 5 lanes of traffic merge into 1




Four Sources of Packet Delay

1. Router processing delay:

- Check bit errors
- Determine output link
- Place packet in buffer

2. Queueing delay

- Time waiting at output link for transmission
- Depends on congestion level of router



Delay in Packet-Switched Networks


3. Transmission delay:

- R = link rate (bps)
- L = packet length (bits)
- Time to send bits into link = L/R


4. Propagation delay:

- d = length of link (m)
- s = propagation speed in medium ($\approx 2 \times 10^8$ m/sec)
- Propagation delay = d/s

Note: s and R are very different quantities!




Caravan Analogy



- Car ~ bit; caravan ~ packet
- Cars “propagate” at 100 mph
- Toll booth takes 12 sec to service a car (transmission time)
- **Q: How long until caravan is lined up before the 2nd toll booth?**
- Time to “push” entire caravan through toll booth onto highway = $12 \times 10 = 120$ sec
- Time for last car to propagate from 1st to 2nd toll booth: $100 \text{ miles} / (100 \text{ mph}) = 1 \text{ hr}$
- **A: 62 minutes**

Caravan Analogy (more)



- Toll booth now takes 1 min to service a car
- **Q: Will cars arrive to 2nd booth before all cars are serviced at 1st booth?**

- **Yes!** After 7 min, 1st car at 2nd booth and 3 cars still at 1st booth
- 1st bit of packet can arrive at 2nd router before packet is fully transmitted from 1st router!
- Can a packet be at 3 routers simultaneously?


Nodal (Per-Router) Delay

$$d_{\text{nodal}} = d_{\text{proc}} + d_{\text{queue}} + d_{\text{trans}} + d_{\text{prop}}$$

- d_{proc} = processing delay
 - A few microseconds or less, usually fixed for all packets
- d_{queue} = queuing delay
 - Depends on congestion, randomly varies between packets
- d_{trans} = transmission delay
 - Equals L/R , high for low-speed links, depends on packet size
- d_{prop} = propagation delay
 - A few microseconds to hundreds of msecs, depends on physical length of the link

Queueing Delay (Revisited)

- R = link bandwidth (bps)
- L = packet length (bits)
- a = average packet arrival rate (pkts/sec)
- Infinite buffer space




Traffic intensity $\rho = La/R$

- $\rho \approx 0$: average queueing delay is small
- $\rho \geq 1$: more “work” arriving than can be serviced, average delay is infinite
- $\rho \rightarrow 1$: delay quickly shoots up

“Real” Internet Delays and Routes


- What do “real” Internet delay & loss look like?
- Traceroute (tracert in Windows): provides delay measurement from source to all routers along end-end Internet path towards destination. For all i :
 - Sends three packets that reach router i on path towards destination
 - Router i returns a response to sender
 - Sender times interval between transmission and reply



“Real” Internet Delays and Routes


traceroute: gaia.cs.umass.edu to www.eurecom.fr

Three delay measurements
at first hop



```
1  cs-gw (128.119.240.254)  1 ms  1 ms  2 ms
2  border1-rt-fa5-1-0.gw.umass.edu (128.119.3.145)  1 ms  1 ms  2 ms
3  cht-vbns.gw.umass.edu (128.119.3.130)  6 ms  5 ms  5 ms
4  jnl-at1-0-0-19.wor.vbns.net (204.147.132.129)  16 ms  11 ms  13 ms
5  jnl-so7-0-0-0.wae.vbns.net (204.147.136.136)  21 ms  18 ms  18 ms
6  abilene-vbns.abilene.ucaid.edu (198.32.11.9)  22 ms  18 ms  22 ms
7  nycm-wash.abilene.ucaid.edu (198.32.8.46)  22 ms  22 ms  22 ms
8  62.40.103.253 (62.40.103.253)  104 ms  109 ms  106 ms
9  de2-1.del.de.geant.net (62.40.96.129)  109 ms  102 ms  104 ms
10 de.fr1.fr.geant.net (62.40.96.50)  113 ms  121 ms  114 ms
11 renater-gw.fr1.fr.geant.net (62.40.103.54)  112 ms  114 ms  112 ms
12 nio-n2.cssi.renater.fr (193.51.206.13)  111 ms  114 ms  116 ms
13 nice.cssi.renater.fr (195.220.98.102)  123 ms  125 ms  124 ms
14 r3t2-nice.cssi.renater.fr (195.220.98.110)  126 ms  126 ms  124 ms
15 eurecom-valbonne.r3t2.ft.net (193.48.50.54)  135 ms  128 ms  133 ms
16 194.214.211.25 (194.214.211.25)  126 ms  128 ms  126 ms
17 * * *
18 * * *
19 fantasia.eurecom.fr (193.55.113.142)  132 ms  128 ms  136 ms
```

* means no reponse (probe lost, router not replying)



Packet Loss

- Queues have **finite** capacity
- When packets arrive to a full buffer, they are dropped (aka lost) – **drop-tail queuing**
- Lost packet may be **retransmitted** by previous router, by the source (end system), or not at all
- **Loss rate**: average fraction of data lost over a long period of time
- Example: link capacity $R = 10$ Mbps and total arrival rate of traffic is 11 Mbps
 - Q: What's the average loss rate on the link?
 - A: About 9%

Chapter 1: Roadmap

- 1.1 What *is* the Internet?
- 1.2 Network edge
- 1.3 Network core
- 1.4 Network access and physical media
- 1.5 Internet structure and ISPs
- 1.6 Delay & loss in packet-switched networks
- 1.7 Protocol layers, service models
- 1.8 History

Protocol “Layers”

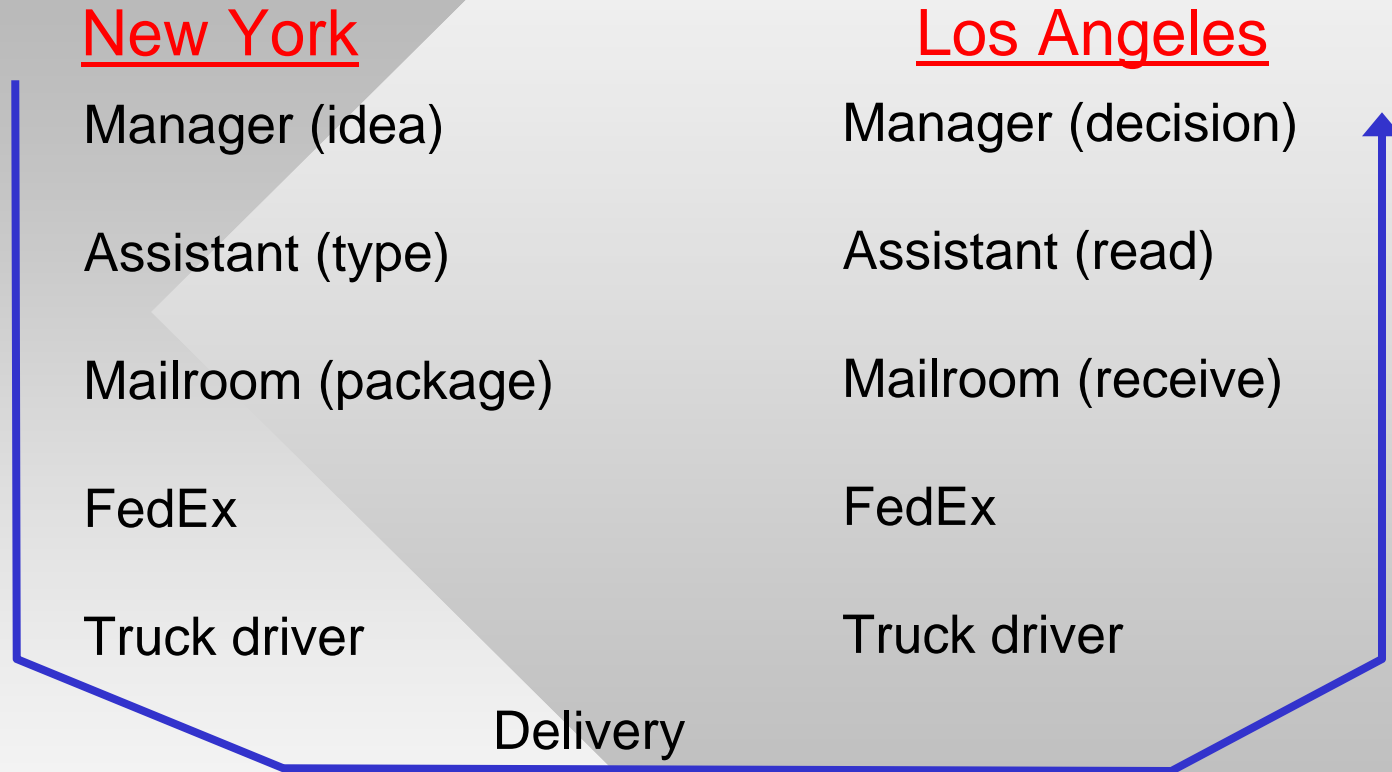
Networks are complex!

- Many “pieces”
 - Hosts
 - Routers
 - Links of various media
 - Applications
 - Protocols
- Some type of modular organization is desirable

Solution: Layered structure


- Same host: each layer interacts only with adjacent (upper/lower) layers
- Remote host: each layer talks to identical layer on the other end-host

Layering



- Information travels **down** the protocol stack on the sender side and **up** on the receiver side

Layering



Layers: each layer implements a service

- Via its own internal-layer actions
- Relying on services provided by the layer below
- Talks to same layer on the other host

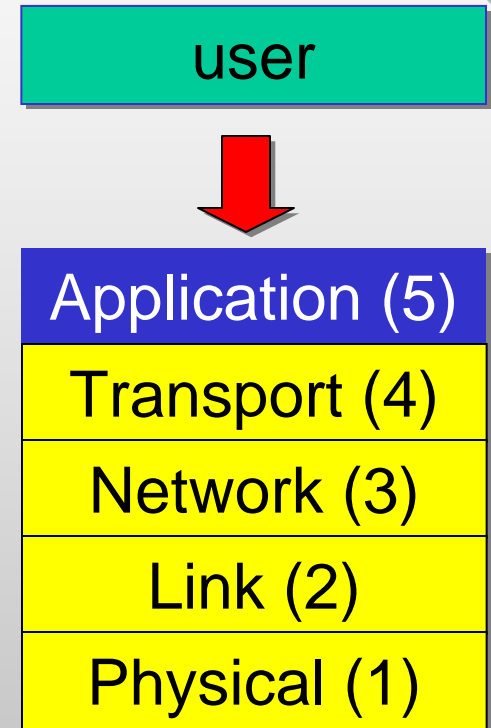
Why Layering?

Benefits of layered organization:

- Sufficient to specify only the **relationship** between the system's pieces
 - Instead of defining one big protocol that does everything
 - Complexity reduced by **separately** standardizing individual components
- Modularization eases maintenance and upgrade
 - Change of implementation of layer's service transparent to the rest of system
 - For example, change in FedEx truck routing doesn't affect other layers







Internet Protocol Stack

- **Application:** interacts with user and supports network applications
 - FTP, SMTP, HTTP (ch 2)
- **Transport:** inter-process data transfer
 - TCP, UDP (ch 3)
- **Network:** routing of datagrams from source to destination host
 - IP, routing protocols (ch 4)
- **Link:** data transfer between neighboring network elements
 - 802.11b, Ethernet (ch 5)
- **Physical:** bits “on the wire”
 - Not covered in this class









Encapsulation

source



switch

destination



router