- Hosts connect to Internet via access Internet Service Providers (ISPs)
  - residential, enterprise (company, university, commercial) ISPs
- Access ISPs in turn must be interconnected
  - so that any two hosts can send packets to each other
- Resulting network of networks is very complex
  - evolution was driven by economics and national policies
- Let's take a stepwise approach to describe current Internet structure

*Question:* given *millions* of access ISPs, how to connect them together?



*Question:* given *millions* of access ISPs, how to connect them together?



*Option:* connect each access ISP to one global transit ISP? *Customer* and *provider* ISPs have economic agreement.



But if one global ISP is viable business, there will be competitors ....



But if one global ISP is viable business, there will be competitors .... who will want to be connected



... and regional networks may arise to connect access nets to ISPs



... and content provider networks (e.g., Google, Microsoft, Akamai) may run their own network, to bring services, content close to end users





At "center": small # of well-connected large networks

- "tier-1" commercial ISPs (e.g., Level 3, Sprint, AT&T, NTT), national & international coverage
- content provider networks (e.g., Google, Facebook): private network that connects its data centers to Internet, often bypassing tier-1, regional ISPs

### Tier-1 ISP Network map: Sprint (2019)



## Chapter 1: roadmap

- What is the Internet?
- What is a protocol?
- Network edge: hosts, access network, physical media
- Network core: packet/circuit switching, internet structure
- Performance: loss, delay, throughput
- Security
- Protocol layers, service models
- History



## How do packet loss and delay occur?

packets queue in router buffers

- packets queue, wait for turn
- arrival rate to link (temporarily) exceeds output link capacity: packet loss



## Packet delay: four sources



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

### *d*<sub>proc</sub>: nodal processing

- check bit errors
- determine output link
- typically < msec</p>

#### d<sub>queue</sub>: queueing delay

- time waiting at output link for transmission
- depends on congestion level of router

## Packet delay: four sources



 $d_{\rm trans}$  and  $d_{\rm prop}$ 

very different

- *d*<sub>trans</sub>: transmission delay:
- L: packet length (bits)

 $d_{trans} = L/R$ 

R: link transmission rate (bps)

#### $d_{\text{prop}}$ : propagation delay:

d: length of physical link

•  $d_{\text{prop}} = d/s$ 

s: propagation speed (~2x10<sup>8</sup> m/sec)

## Caravan analogy



- cars "propagate" at 100 km/hr
- toll booth takes 12 sec to service car (bit transmission time)
- car ~ bit; caravan ~ packet
- Q: How long until caravan is lined up before 2nd toll booth?

- time to "push" entire caravan through toll booth onto highway = 12\*10 = 120 sec
- time for last car to propagate from 1st to 2nd toll both: 100km/(100km/hr) = 1 hr
- A: 62 minutes

## Caravan analogy



- suppose cars now "propagate" at 1000 km/hr
- and suppose toll booth now takes one min to service a car
- Q: Will cars arrive to 2nd booth before all cars serviced at first booth?
  <u>A: Yes!</u> after 7 min, first car arrives at second booth; three cars still at first booth

## Packet queueing delay (revisited)

- R: link bandwidth (bps)
- L: packet length (bits)
- *a:* average packet arrival rate
- La/R ~ 0: avg. queueing delay small
- La/R -> 1: avg. queueing delay large
- La/R > 1: more "work" arriving is more than can be serviced - average delay infinite!





## "Real" Internet delays and routes

- what do "real" Internet delay & loss look like?
- traceroute program: provides delay measurement from source to router along end-end Internet path towards destination. For all *i*:
  - sends three packets that will reach router *i* on path towards destination (with time-to-live field value of *i*)
  - router *i* will return packets to sender
  - sender measures time interval between transmission and reply



## **Real Internet delays and routes**

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

3 delay measurements from gaia.cs.umass.edu to cs-gw.cs.umass.edu 1 cs-gw (128.119.240.254) 1 ms 1 ms 2 ms to border1-rt-fa5-1-0.gw.umass.edu 3 cht-vbns.gw.umass.edu (128.119.3.130) 6 ms 5 ms 5 ms 4 jn1-at1-0-0-19.wor.vbns.net (204.147.132.129) 16 ms 11 ms 13 ms 5 jn1-so7-0-0.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 1 trans-oceanic link 8 62.40.103.253 (62.40.103.253) 104 ms 109 ms 106 ms 🛶 9 de2-1.de1.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 looks like delays 11 renater-gw.fr1.fr.geant.net (62.40.103.54) 112 ms 114 ms 112 ms decrease! Why? 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 \* \* \* \* means no response (probe lost, router not replying) \* \* \* 18 19 fantasia.eurecom.fr (193.55.113.142) 132 ms 128 ms 136 ms

\* Do some traceroutes from exotic countries at www.traceroute.org

### Packet loss

- queue (aka buffer) preceding link in buffer has finite capacity
- packet arriving to full queue dropped (aka lost)
- lost packet may be retransmitted by previous node, by source end system, or not at all



# Throughput

- throughput: rate (bits/time unit) at which bits are being sent from sender to receiver
  - *instantaneous:* rate at given point in time
  - *average:* rate over longer period of time



## Throughput

