

# Computer Networking

## Slide Set 1

**Andrew W. Moore**

Andrew.Moore@cl.cam.ac.uk

## Topic 1 Foundation

- Administrivia
- Networks
- Channels
- Multiplexing
- Performance: loss, delay, throughput

3

## Course Administration

### Commonly Available Texts

- ❑ Computer Networks: A Systems Approach  
Peterson and Davie  
<https://book.systemsapproach.org>  
<https://github.com/SystemsApproach/book>

### Other Selected Texts (non-representative)


- ❑ Computer Networking: A Top-Down Approach  
Kurose and Ross, (many editions), Addison-Wesley
- ❑ Internetworking with TCP/IP, vol. I + II  
Comer & Stevens, Prentice Hall
- ❑ UNIX Network Programming, Vol. I  
Stevens, Fenner & Rudoff, Prentice Hall



4

## What is a network?

- A system of “links” that interconnect “nodes” in order to move “information” between nodes




- Yes, this is all rather abstract


6

## What is a network?

- We also talk about



or even



- Yes, abstract, vague, and under-defined....

7

## There are *many* different types of networks

- Internet
- Telephone network
- Transportation networks
- Cellular networks
- Supervisory control and data acquisition networks
- Optical networks
- Sensor networks

**We will focus almost exclusively on the Internet**

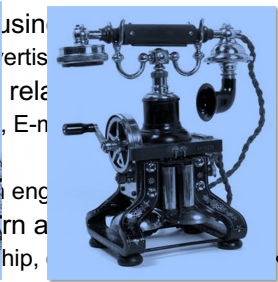
8

## The Internet has transformed everything

- The way we do business
  - E-commerce, advertising, cloud-computing
- The way we have relationships
  - Facebook friends, E-mail, IM, virtual worlds
- The way we learn
  - Wikipedia, search engines
- The way we govern and view law
  - E-voting, censorship, copyright, cyber-attacks

9

## The Internet transforms everything



Taking the dissemination of information to the next level

## The Internet is big business

- Many large and influential networking companies
  - Huawei, Broadcom, AT&T, Verizon, Akamai, Cisco, ...
  - \$132B+ industry (carrier and enterprise alone)
- Networking central to most technology companies
  - Apple, Google, Facebook, Intel, Amazon, VMware, ...


## But why is the Internet *interesting*?

- “What’s your formal model for the Internet?” -- *theorists*
- “Aren’t you just writing software for networks” – *hackers*
- “You don’t have performance benchmarks???” – *hardware folks*
- “Isn’t it just another network?” – *old timers at BT*
- “What’s with all these TLA protocols?” – *all*
- “But the Internet seems to be working...” – *my mother*

## A few defining characteristics of the Internet

## A federated system

- The Internet ties together different networks
  - >20,000 ISP networks (the definition is fuzzy)



Tied together by IP -- the “Internet Protocol” : a single common interface between users and the network and between networks

## A federated system

- The Internet ties together different networks
  - >20,000 ISP networks
- A single, common interface is great for interoperability...
- ...but tricky for business
- Why does this matter?
  - ease of interoperability is the Internet’s most important goal
  - practical realities of incentives, economics and real-world trust, drive topology, route selection and service evolution

## Tremendous scale (2020 numbers – so some ‘weird’)

- 4.57 Billion users (58% of world population)
- 1.8 Billion web sites
  - 34.5% of which are powered by the WordPress!
- 4.88 Billion smartphones (45.4% of population)
- 500 Million Tweets a day
- 100 Billion WhatsApp messages per day
- 1 Billion hours of YouTube video watched per day
- 500 hours of Youtube video added per minute
- 2+ billion TikTok installs
- 60% video streaming
  - 12.5% of the Internet traffic is native Netflix

17

## Tremendous scale (2020 numbers – so some ‘weird’)

- 4.57 Billion users (58% of world population)
- 1.8 Billion web sites
  - 34.5% of which are powered by the WordPress!
- 4.88 Billion smartphones (45.4% of population)
- 500 Million Tweets a day
- 100 Billion WhatsApp messages per day
- 1 Billion hours of YouTube video watched per day
- 500 hours of Youtube video added per minute
- 2+ billion TikTok installs
- 60% video streaming
  - 12.5% of the Internet traffic is native Netflix

18

## Enormous diversity and dynamic range

- Communication latency: microseconds to seconds ( $10^6$ )
- Bandwidth: 1Kbits/second to 400 Gigabits/second ( $10^7$ )
- Packet loss: 0 – 90%
- Technology: optical, wireless, satellite, copper
- **Endpoint devices:** from sensors and cell phones to datacenters and supercomputers
- **Applications:** social networking, file transfer, skype, live TV, gaming, remote medicine, backup, IM
- **Users:** the governing, governed, operators, malicious, naïve, savvy, embarrassed, paranoid, addicted, cheap ...

19

## Constant Evolution

1970s:

- 56kilobits/second “backbone” links
- <100 computers, a handful of sites in the US (and one UK)
- Telnet and file transfer are the “killer” applications

Today

- 400+Gigabits/second backbone links
- 40B+ devices, all over the globe
  - 27B+ IoT devices alone

20

## Asynchronous Operation

- Fundamental constraint: **speed of light**
- Consider:
  - How many cycles does your 3GHz CPU in Cambridge execute before it can possibly get a response from a message it sends to a server in Palo Alto?
    - Cambridge to Palo Alto: 8,609 km
    - Traveling at 300,000 km/s: 28.70 milliseconds
    - Then back to Cambridge:  $2 \times 28.70 = 57.39$  milliseconds
    - $3,000,000,000 \text{ cycles/sec} \times 0.05739 = 172,179,999$  cycles!
- Thus, communication feedback is always *dated*

21

## Prone to Failure

- To send a message, **all** components along a path must function correctly
  - software, wireless access point, firewall, links, network interface cards, switches, ...
  - Including **human operators**
- Consider: 50 components, that work correctly 99% of time → 39.5% chance communication will fail
- Plus, recall
  - scale → lots of components
  - asynchrony → takes a long time to hear (bad) news
  - federation (**internet**) → hard to identify fault or assign blame

22

## Recap: The Internet is...

- A complex federation
- Of enormous scale
- Dynamic range
- Diversity
- Constantly evolving
- Asynchronous in operation
- Failure prone
- Constrained by what's practical to engineer
- Too complex for (simple) theoretical models
- "Working code" doesn't mean much
- Performance benchmarks are too narrow


23

## An Engineered System

- Constrained by what technology is practical
  - Link bandwidths
  - Switch port counts
  - Bit error rates
  - Cost
  - ...

24


## Nodes and Links



Channels = Links  
Peer entities = Nodes

25

## Properties of Links (Channels)



- Bandwidth (capacity): "width" of the links
  - number of bits sent (or received) per unit time (bits/sec or bps)
- Latency (delay): "length" of the link
  - propagation time for data to travel along the link (seconds)
- Bandwidth-Delay Product (BDP): "volume" of the link
  - amount of data that can be "in flight" at any time
  - propagation delay  $\times$  bits/time = total bits in link

26

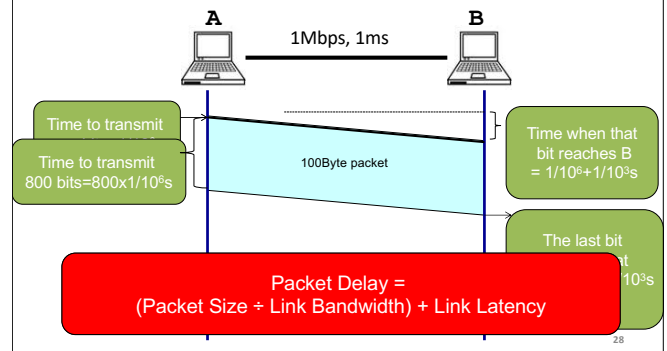
## Examples of Bandwidth-Delay

- Same city over a slow link:
  - BW~100Mbps
  - Latency~10msec
  - BDP ~  $10^6$ bits ~ 125KBytes
- Intra Datacenter:
  - BW~100Gbps
  - Latency~30usec
  - BDP ~  $10^6$ bits ~ 375KBytes
- Cross-Atlantic over fast link:
  - BW~10Gbps
  - Latency~100msec
  - BDP ~  $10^9$ bits ~ 125MBytes
- Intra Host:
  - BW~100Gbps
  - Latency~16nsec
  - BDP ~ 1600bits ~ 200Bytes

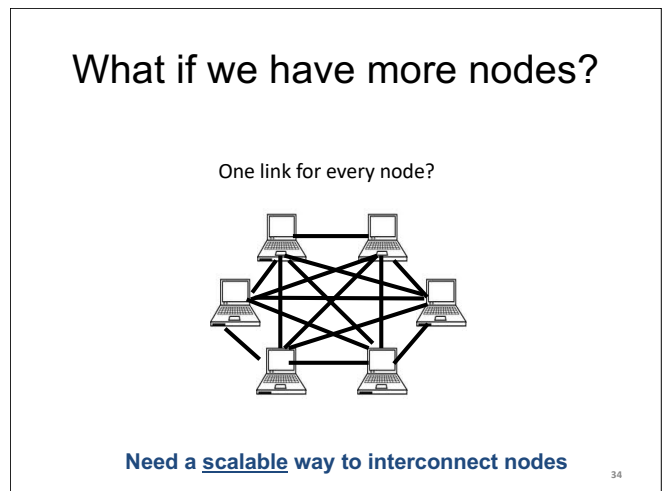
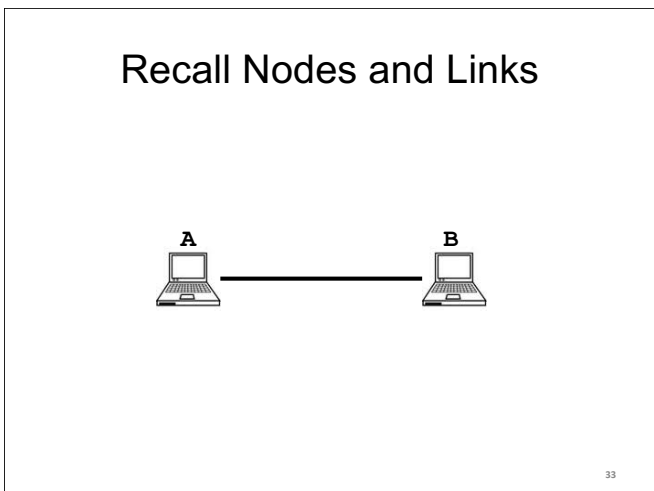
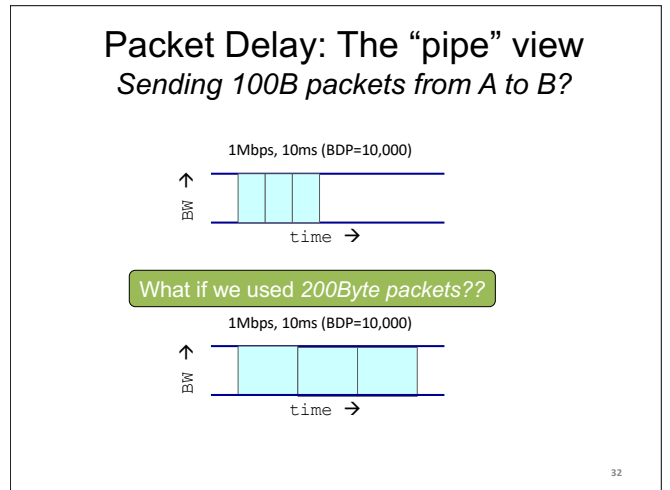
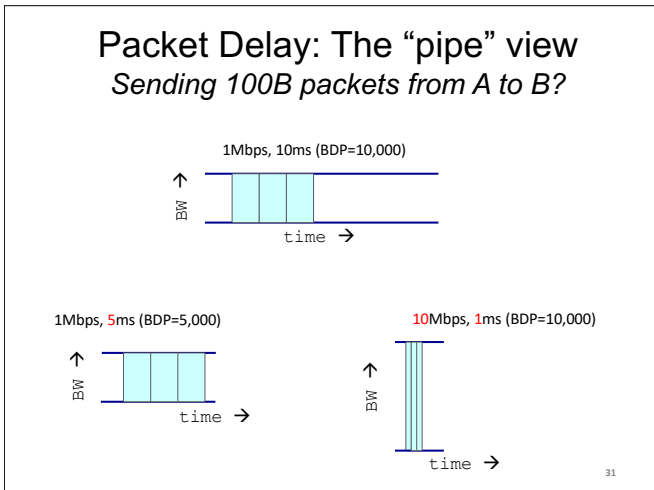
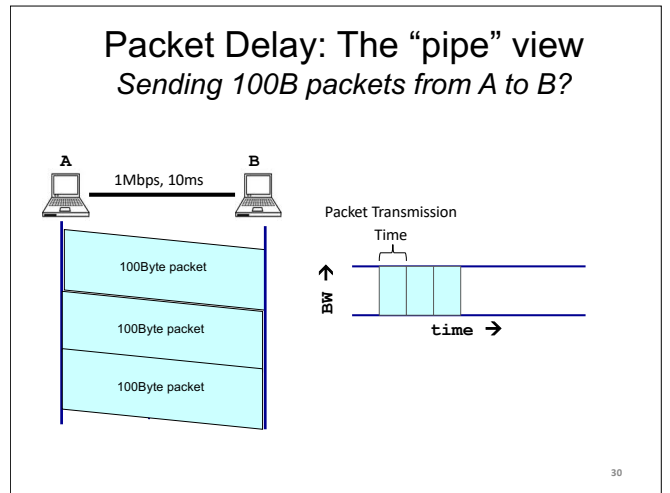
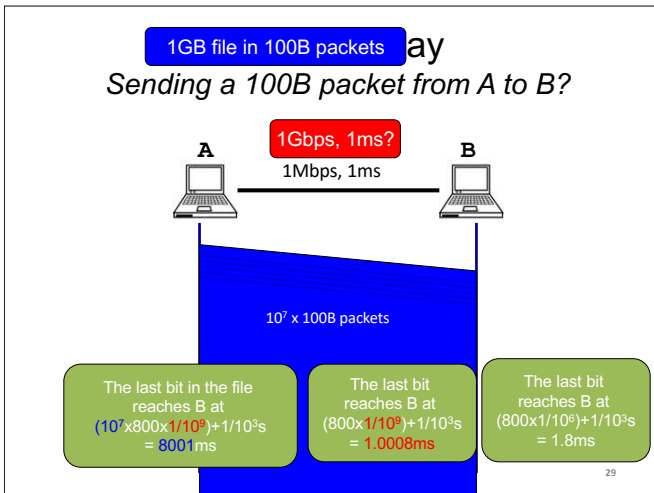
27

## Packet Delay

Sending a 100B packet from A to B?




28



## Solution: A switched network

Nodes share network link resources



How is this sharing implemented?

35


## Two forms of switched networks

- Circuit switching (used in the *POTS*: Plain Old Telephone system)
- Packet switching (used in the Internet)

36

## Circuit switching

Idea: source **reserves** network capacity along a path



- (1) Node A sends a reservation request
- (2) Interior switches establish a connection -- i.e., "circuit"
- (3) A starts sending data
- (4) A sends a "teardown circuit" message

37

## Multiplexing



Sharing makes things efficient (cost less)

- One airplane/train for 100's of people
- One telephone for many calls
- One lecture theatre for many classes
- One computer for many tasks
- One network for many computers
- One datacenter many applications

38

## Multiplexing




Sharing makes things efficient (cost less)

- One airplane/train for 100's of people
- One telephone for many calls
- One ~~lecture theatre~~ **Lecturer?** for many classes
- One computer for many tasks
- One network for many computers
- One datacenter many applications

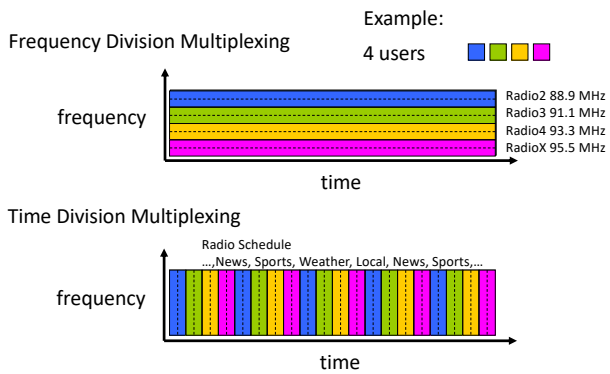
39

## Old Time Multiplexing




40

## Circuit Switching: FDM and TDM




## Time-Division Multiplexing/Demultiplexing



- Time divided into frames; frames into slots
- Relative slot position inside a frame determines to which conversation data belongs
  - e.g., slot 0 belongs to orange conversation
- Slots are reserved (released) during circuit setup (teardown)
- If a conversation does not use its circuit **capacity is lost!**


## Timing in Circuit Switching



## Circuit switching: pros and cons

- Pros
  - guaranteed performance
  - fast transfer (once circuit is established)
- Cons

## Timing in Circuit Switching




## Circuit switching: pros and cons

- Pros
  - guaranteed performance
  - fast transfer (once circuit is established)
- Cons
  - **wastes bandwidth if traffic is “bursty”**




## Timing in Circuit Switching



47

## Timing in Circuit Switching




48

## Circuit switching: pros and cons

- Pros
  - guaranteed performance
  - fast transfers (once circuit is established)
- Cons
  - wastes bandwidth if traffic is “bursty”
  - **connection setup time is overhead**

49

## Circuit switching



Circuit switching doesn't "route around failure"

50

## Circuit switching: pros and cons

- Pros
  - guaranteed performance
  - fast transfers (once circuit is established)
- Cons
  - wastes bandwidth if traffic is “bursty”
  - connection setup time is overhead
  - **recovery from failure is slow**

51

## Numerical example

- How long does it take to send a file of 640,000 bits from host A to host B over a circuit-switched network?
  - All links are 1.536 Mbps
  - Each link uses TDM with 24 slots/sec
  - 500 msec to establish end-to-end circuit

Let's work it out!

52


## Two forms of switched networks

- Circuit switching (e.g., telephone network)
- Packet switching (e.g., Internet)

53

## Packet Switching

- Data is sent as chunks of formatted bits (**Packets**)
- Packets consist of a “**header**” and “**payload**”\*



55

## Packet Switching

- Data is sent as chunks of formatted bits (**Packets**)
- Packets consist of a “**header**” and “**payload**”\*
  - payload is the data being carried
  - header holds instructions to the network for how to handle packet (think of the header as an API)
- In this example, the header has a destination address
- More complex headers may include
  - How this traffic should be handled? (first class, second class, etc)
  - Who signed for it?
  - Were the contents ok?


56

## Packet Switching

- Data is sent as chunks of formatted bits (Packets)
- Packets consist of a “header” and “payload”
- Switches “**forward**” packets based on their headers

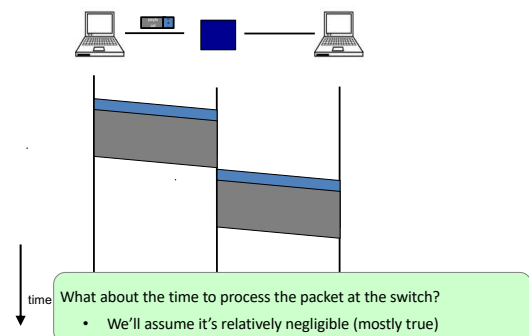
57

## Switches forward packets




58

## Timing in Packet Switching



59


## Timing in Packet Switching



Could the switch start transmitting as soon as it has processed the header?

- Yes! This would be called a "cut through" switch

## Timing in Packet Switching



We will always assume a switch processes/forwards a packet after it has received it entirely. This is called "store and forward" switching

## Packet Switching

- Data is sent as chunks of formatted bits (Packets)
- Packets consist of a "header" and "payload"
- Switches "forward" packets based on their headers

## Packet Switching

- Data is sent as chunks of formatted bits (Packets)
- Packets consist of a "header" and "payload"
- Switches "forward" packets based on their headers
- Each packet travels independently
  - no notion of packets belonging to a "circuit"

## Packet Switching

- Data is sent as chunks of formatted bits (Packets)
- Packets consist of a "header" and "payload"
- Switches "forward" packets based on their headers
- Each packet travels independently
- No link resources are reserved in advance. Instead packet switching leverages **statistical multiplexing** (stat muxing)


## Multiplexing



Sharing makes things efficient (cost less)


- One airplane/train for 100's of people
- One telephone for many calls
- One lecture theatre for many classes
- One computer for many tasks
- One network for many computers
- One datacenter many applications

## Three Flows with Bursty Traffic



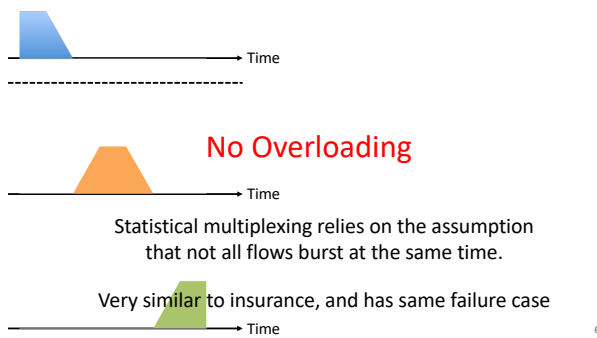
66

## When Each Flow Gets 1/3<sup>rd</sup> of Capacity




67

## When Flows Share Total Capacity



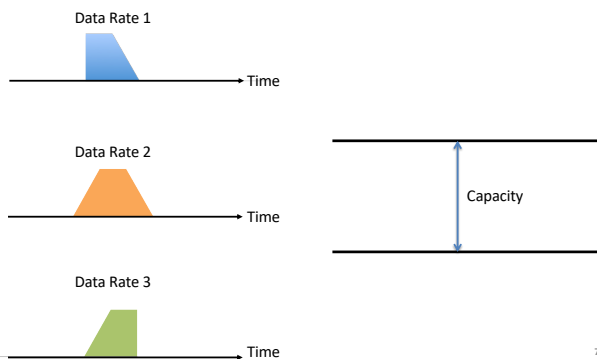
68

## Three Flows with Bursty Traffic




69

## Three Flows with Bursty Traffic



70

## Three Flows with Bursty Traffic




What do we do under overload?

71

Sorry we don't carry https here....




Statistical multiplexing: pipe view




73

Statistical multiplexing: pipe view




74

Statistical multiplexing: pipe view




75

Statistical multiplexing: pipe view




76

Statistical multiplexing: pipe view




77

## Statistical multiplexing: pipe view




78

## Statistical multiplexing: pipe view



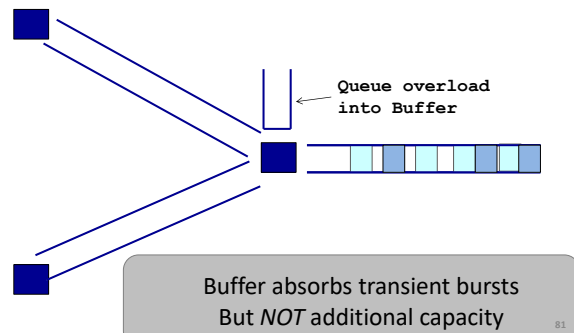
79

## Statistical multiplexing: pipe view



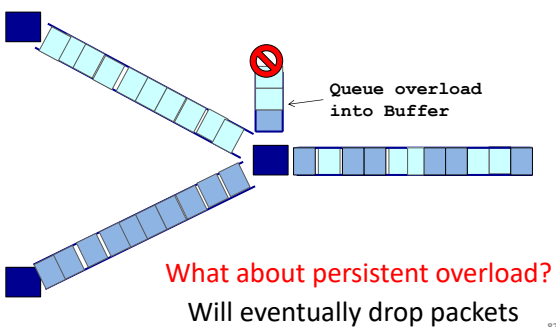
80

## Statistical multiplexing: pipe view



81

## Statistical multiplexing: pipe view



82

## Queues introduce queuing delays


- Recall,
 
$$\text{packet delay} = \text{transmission delay} + \text{propagation delay} (*)$$
- With queues (statistical multiplexing)
 
$$\text{packet delay} = \text{transmission delay} + \text{propagation delay} + \text{queuing delay} (*)$$
- Queuing delay caused by "packet interference"
- Made worse at high load
  - less "idle time" to absorb bursts
  - think about traffic jams at rush hour or rail network failure

(\* plus per-hop *processing* delay that we define as negligible)

83

## Queuing delay extremes

- $R$ =link bandwidth (bps)
- $L$ =packet length (bits)
- $a$ =average packet arrival rate




traffic intensity =  $La/R$

- $La/R \sim 0$ : average queuing delay small
- $La/R \rightarrow 1$ : delays become large
- $La/R > 1$ : more "work" arriving than can be serviced, average delay infinite – or data is lost (*dropped*).

84

## Recall the Internet *federation*

- The Internet ties together different networks
  - >20,000 ISP networks



We can see (hints) of the nodes and links using traceroute...

85

## "Real" Internet delays and routes

traceroute: rio.cl.cam.ac.uk to munnari.oz.au

(tracepath on windows is similar)

Three delay measurements from

rio.cl.cam.ac.uk to gatwick.net.cl.cam.ac.uk

trans-continent link


```

traceroute munnari.oz.au
traceroute to munnari.oz.au (202.29.151.3), 30 hops max, 60 byte packets
 1  gatwick.net.cl.cam.ac.uk (128.232.32.2)  0.416 ms  0.384 ms  0.427 ms
 2  cl-sby.route-nwest.net.cam.ac.uk (193.60.89.9)  0.393 ms  0.440 ms  0.494 ms
 3  route-nwest.route-mill.net.cam.ac.uk (192.84.5.137)  0.407 ms  0.448 ms  0.501 ms
 4  route-mill.route-enet.net.cam.ac.uk (192.84.5.94)  1.006 ms  1.091 ms  1.163 ms
 5  xe-11-3-0.cambridge-eastern.ja.net (146.97.130.1)  0.300 ms  0.313 ms  0.350 ms
 6  ae24.londix-sbrl.ja.net (146.97.37.185)  2.679 ms  2.664 ms  2.712 ms
 7  ae28.londix-sbrl.ja.net (146.97.33.17)  5.955 ms  5.953 ms  5.901 ms
 8  janet.mx1.lon.uk.geant.net (62.40.124.197)  6.059 ms  6.066 ms  6.052 ms
 9  ae0.mx1.par.fr.geant.net (62.40.98.77)  11.742 ms  11.779 ms  11.724 ms
10  ae1.mx1.mad.es.geant.net (62.40.98.64)  27.751 ms  27.734 ms  27.704 ms
11  mb-so-02-v4.bb.tein3.net (202.179.249.117)  138.296 ms  138.314 ms  138.282 ms
12  sg-so-04-v4.bb.tein3.net (202.179.249.53)  196.303 ms  196.293 ms  196.264 ms
13  th-pr-v4.bb.tein3.net (202.179.249.66)  225.153 ms  225.178 ms  225.196 ms
14  pyt-thairen-to-02-bdr-pyt.uni.net.th (202.29.12.10)  225.163 ms  223.343 ms  223.363 ms
15  202.28.227.126 (202.28.227.126)  241.038 ms  240.941 ms  240.834 ms
16  202.28.221.46 (202.28.221.46)  287.252 ms  287.306 ms  287.282 ms
17  *
18  *
19  *
20  coe-gw.psu.ac.th (202.29.149.70)  241.681 ms  241.715 ms  241.680 ms
21  munnari.OZ.AU (202.29.151.3)  241.610 ms  241.636 ms  241.537 ms
    
```

86

## Internet structure: network of networks


- a packet passes through many networks!



87

## Internet structure: network of networks


- "Tier-3" ISPs and local ISPs
  - last hop ("access") network (closest to end systems)



88

## Internet structure: network of networks


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



89

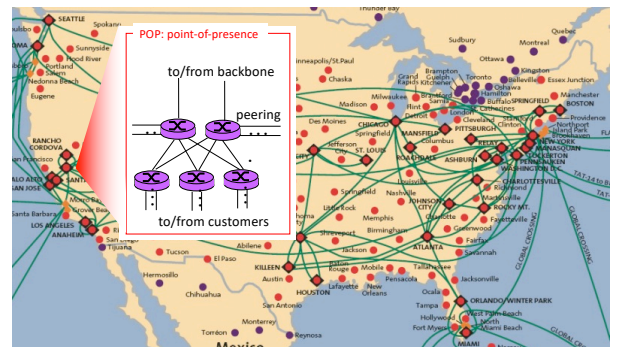
## Internet structure: network of networks

- roughly hierarchical
- at center: "tier-1" ISPs (e.g., Verizon, Sprint, AT&T, Cable and Wireless), national/international coverage
  - treat each other as equals



90

## Tier-1 ISP: e.g., Sprint



91

## Packet Switching


- Data is sent as chunks of formatted bits (Packets)
- Packets consist of a "header" and "payload"
- Switches "forward" packets based on their headers
- Each packet travels independently
- No link resources are reserved in advance. Instead packet switching leverages **statistical multiplexing**
  - allows efficient use of resources
  - but introduces queues and queuing delays

92

## Packet switching versus circuit switching

*Packet switching may (does!) allow more users to use network*

- 1 Mb/s link
- each user:
  - 100 kb/s when "active"
  - active 10% of time
- **circuit-switching:**
  - 10 users
- **packet switching:**
  - with 35 users, probability > 10 active at same time is less than .0004



Q: how did we get value 0.0004?

93

## Packet switching versus circuit switching

Q: how did we get value 0.0004?

- 1 Mb/s link
- each user:
  - 100 kb/s when "active"
  - active 10% of time
- **circuit-switching:**
  - 10 users
- **packet switching:**
  - with 35 users, probability > 10 active at same time is less than .0004

Let U be number of users active  
N the total users  
P is 0.1 in our example to get 0.0004

95

$$\hat{P}(u = k) = \binom{n}{k} p^k (1-p)^{n-k}$$

$$\left[ P(u \leq K) = \sum_{k=0}^K \binom{n}{k} p^k (1-p)^{n-k} \right] \left[ P(u > K) = 1 - \sum_{k=0}^K \binom{n}{k} p^k (1-p)^{n-k} \right]$$

for  $n = 35, K = 10$

$$P(u \leq 10) = \sum_{k=0}^{10} \binom{35}{k} p^k (1-p)^{35-k}$$

where  $p = 0.1$ :

$$P(u \leq 10) = 0.99958$$

$$\therefore P(u > 10) = 0.00042$$

96



## Circuit switching: pros and cons

- Pros
  - guaranteed performance
  - fast transfers (once circuit is established)
- Cons
  - wastes bandwidth if traffic is “bursty”
  - connection setup adds delay
  - recovery from failure is slow

98

## Packet switching: pros and cons

- Cons
  - no guaranteed performance
  - header overhead per packet
  - queues and queuing delays
- Pros
  - efficient use of bandwidth (stat. muxing)
  - no overhead due to connection setup
  - resilient -- can `route around trouble`

99

## Summary

- A sense of how the basic `plumbing` works
  - links and switches
  - packet delays= transmission + propagation + queuing + (negligible) per-switch processing
  - statistical multiplexing and queues
  - circuit vs. packet switching

100

## Topic 2 – Architecture and Philosophy


- Abstraction
- Layering
- Layers and Communications
- Entities and Peers
- What is a protocol?
- Protocol Standardization
- The architects process
  - How to break system into modules
  - Where modules are implemented
  - Where is state stored
- Internet Philosophy and Tensions

1

## TRIGGER WARNING

- Philosophy,
- Bad Analogies, and
- RANTS verging on POLEMIC

Will follow.....



## Abstraction Concept

A mechanism for breaking down a problem

*what not how*

- eg Specification *versus* implementation
- eg Modules in programs

Allows replacement of implementations without affecting system behavior

*Vertical versus Horizontal*

*"Vertical"* what happens in a box "How does it attach to the network?"

*"Horizontal"* the communications paths running through the system

**Hint:** paths are built ("layered") on top of other paths

3

## Computer System Modularity

Partition system into modules & abstractions:

- Well-defined interfaces give flexibility
  - **Hides** implementation - can be freely changed
  - Extend functionality of system by adding new modules
- E.g., libraries encapsulating set of functionality
- E.g., programming language + compiler abstracts away how the particular CPU works ...

4

## Computer System Modularity (cnt' d)

- Well-defined interfaces hide information
  - Isolate **assumptions**
  - Present high-level **abstractions**
- **But can impair performance!**
- Ease of implementation vs worse performance

5

## Network System Modularity

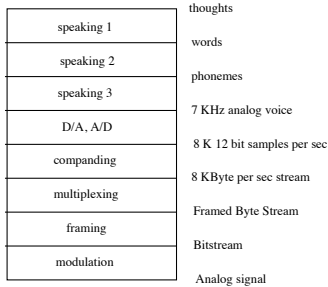
Like software modularity, but:

- Implementation is distributed across many machines (routers and hosts)
- Must decide:
  - How to break system into modules
    - **Layering**
  - Where modules are implemented
    - **End-to-End Principle**
  - Where state is stored
    - **Fate-sharing**

6

## Layering Concept


- A restricted form of abstraction: system functions are divided into layers, one built upon another
- Often called a *stack*; but **not** a data structure!



7

## Layers and Communications

- Interaction only between adjacent layers
- *layer n* uses services provided by *layer n-1*
- *layer n* provides service to *layer n+1*
- Bottom layer is physical media
- Top layer is application



8

## Entities and Peers


*Entity* – a *thing* (an independent existence)

Entities *interact* with the layers above and below

Entities *communicate* with *peer* entities

- same level but different place (eg different person, different box, different host)

Communications between peers is supported by entities at the lower layers



9

## Entities and Peers


Entities usually do something useful

- Encryption – Error correction – Reliable Delivery
- Nothing at all is also reasonable

Not all communications is end-to-end

Examples for things in the middle

- IP Router – Mobile Phone Cell Tower
- Person translating French to English



10

## Layering and Embedding


In Computer Networks we often see higher-layer information embedded within lower-layer information

- Such embedding can be considered a form of layering
- Higher layer information is generated by stripping off headers and trailers of the current layer
- eg an IP entity only looks at the IP headers

**BUT embedding is not the only form of layering**

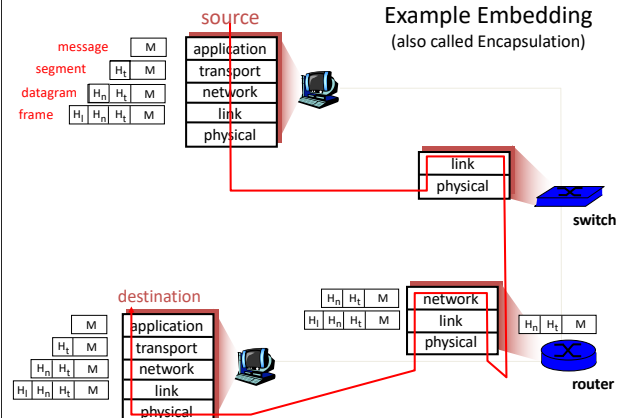
Layering is to help understand a communications system

**NOT** determine implementation strategy



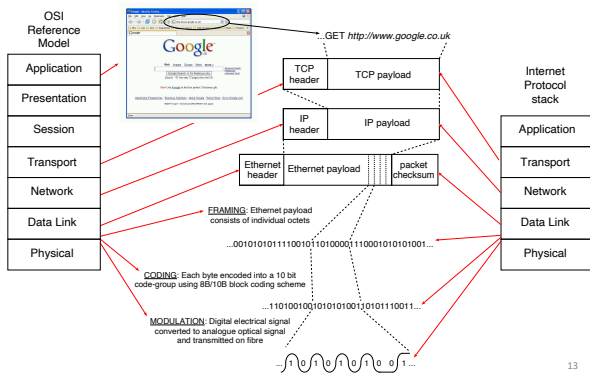
11

## Example Embedding (also called Encapsulation)



12

## Internet protocol stack *versus* OSI Reference Model



13

## ISO/OSI reference model

- **presentation**: allow applications to interpret meaning of data, e.g., encryption, compression, machine-specific conventions
- **session**: synchronization, checkpointing, recovery of data exchange
- Internet stack “missing” these layers!
  - these services, *if needed*, must be implemented in application

application
presentation
session
transport
network
link
physical

14

## What is a protocol?

### human protocols:

- “what’s the time?”
- “I have a question”
- introductions

... specific msgs sent

... specific actions taken when msgs received, or other events

### network protocols:

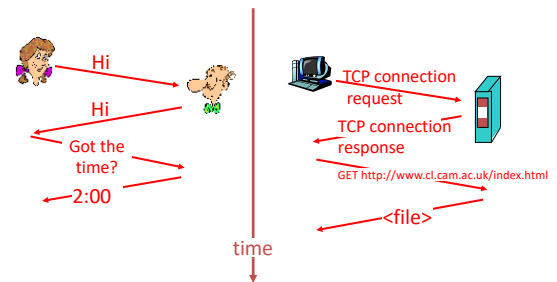
- machines rather than humans
- all communication activity in Internet governed by protocols

*protocols define format, order of msgs sent and received among network entities, and actions taken on msg transmission, receipt*

15

## What is a protocol?

a human protocol and a computer network protocol:



Q: Other human protocols?

16


## Protocol Standardization

- All hosts must follow same protocol
  - Very small modifications can make a big difference
  - Or prevent it from working altogether
- This is why we have standards
  - Can have multiple implementations of protocol
- Internet Engineering Task Force (IETF)
  - Based on working groups that focus on specific issues
  - Produces “Request For Comments” (RFCs)
  - IETF Web site is <http://www.ietf.org>
  - RFCs archived at <http://www.rfc-editor.org>

17


## So many Standards Problem

- Many different packet-switching networks
- Each with its own Protocol
- Only nodes on the same network could communicate



18

## INTERNET Solution



19

## Internet Design Goals (Clark '88)

- **Connect existing networks**
- Robust in face of failures
- Support multiple types of delivery services
- Accommodate a variety of networks
- Allow distributed management
- Easy host attachment
- Cost effective
- Allow resource accountability

20

## Real Goals


Internet Motto

*We reject kings , presidents, and voting. We believe in rough consensus and running code.* – David Clark

- **Build something that works!**
- Connect existing networks
- Robust in face of failures
- Support multiple types of delivery services
- Accommodate a variety of networks
- Allow distributed management
- Easy host attachment
- Cost effective
- ~~Allow resource accountability~~

21

## A Multitude of Apps Problem




- Re-implement every application for every technology?
- No! But how does the Internet design avoid this?

22

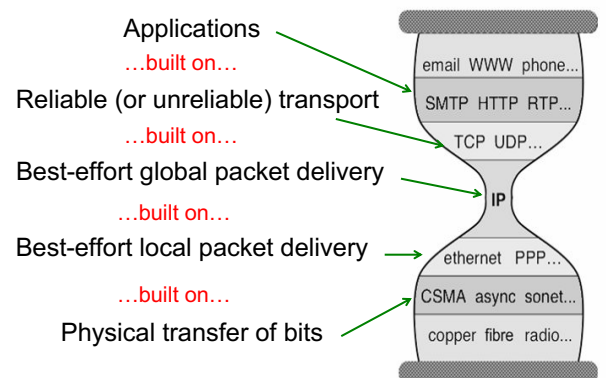
## Solution: Intermediate Layers

- Introduce intermediate layers that provide **set of abstractions** for various network functionality and technologies
  - A new app/media implemented only once
  - Variation on “add another level of indirection”



23


## In the context of the Internet



24

## Three Observations


- Each layer:
  - Depends on layer below
  - Supports layer above
  - Independent of others
- Multiple versions in layer
  - Interfaces differ somewhat
  - Components pick which lower-level protocol to use
- But only one IP layer
  - Unifying protocol



25

## Layering Crucial to Internet's Success

- Reuse
- Hides underlying detail
- Innovation at each level can proceed in parallel
- Pursued by very different communities



26

What are some of the drawbacks of protocols and layering?

27

## Drawbacks of Layering

- Layer N may duplicate lower layer functionality
  - e.g., error recovery to retransmit lost data
- Information hiding may hurt performance
  - e.g., packet loss due to corruption vs. congestion
- Headers start to get really big
  - e.g., typical TCP+IP+Ethernet is 54 bytes
- Layer violations when the gains too great to resist
  - e.g., TCP-over-wireless
- Layer violations when network doesn't trust ends
  - e.g., firewalls

28

## Placing Network Functionality

- Hugely influential paper: "End-to-End Arguments in System Design" by Saltzer, Reed, and Clark ('84)
  - articulated as the "End-to-End Principle" (E2E)
- Endless debate over what it means
- Everyone cites it as supporting their position  
(regardless of the position!)


29

## Basic Observation

- Some application requirements can only be correctly implemented **end-to-end**
  - reliability, security, etc.
- Implementing these in the network is hard
  - every step along the way must be fail proof
- Hosts
  - **Can** satisfy the requirement without network's help
  - **Will/must** do so, since they can't rely on the network

30


## Example: Reliable File Transfer



- Solution 1: make each step reliable, and string them together to make reliable end-to-end process

31

## Example: Reliable File Transfer



- Solution 1: make each step reliable, and string them together to make reliable end-to-end process


So what is the problem?

each component is 0.9 reliable

leads to total system failure of  $>0.4^*$

32

## Example: Reliable File Transfer



- Solution 1: make each step reliable, and string them together to make reliable end-to-end process
- Solution 2: end-to-end **check** and retry

33

## Discussion

- Solution 1 is incomplete
  - What happens if any network element misbehaves?
  - Receiver has to do the check anyway!
- Solution 2 is complete
  - Full functionality can be entirely implemented at application layer with no need for reliability from lower layers
- Is there any need to implement reliability at lower layers?

34

## Summary of End-to-End Principle

- Implementing functionality (e.g., reliability) in the network
  - Doesn't reduce host implementation complexity
  - Does increase network complexity
  - Probably increases delay and overhead on all applications even if they don't need the functionality (e.g. VoIP)
- However, implementing in the network can improve performance in some cases
  - e.g., consider a very lossy link

35

## “Only-if-Sufficient” Interpretation

- Don't implement a function at the lower levels of the system unless it can be completely implemented at this level
- *Unless you can relieve the burden from hosts, don't bother*

36

## “Only-if-Necessary” Interpretation

- Don't implement *anything* in the network that can be implemented correctly by the hosts
- Make network layer absolutely minimal
  - This E2E interpretation trumps performance issues
  - Increases flexibility, since lower layers stay **simple**

37

## “Only-if-Useful” Interpretation

- If hosts can implement functionality correctly, implement it in a lower layer **only** as a performance enhancement
- But do so only if it **does not impose burden** on applications that do not require that functionality

38

## We have some tools:

- Abstraction
- Layering
- Layers and Communications
- Entities and Peers
- Protocol as motivation
- Examples of the architects process
- Internet Philosophy and Tensions

39


## Distributing Layers Across Network

- Layers are simple if only on a single machine
  - Just stack of modules interacting with those above/below
- But we need to implement layers across machines
  - Hosts
  - Routers (switches)
- What gets implemented where?

40

## What Gets Implemented on Host?


- Bits arrive on wire, must make it up to application
- Therefore, all layers must exist at the host



41

## What Gets Implemented on a Router?

- Bits arrive on wire
  - Physical layer necessary
- Packets must be delivered to next-hop
  - Datalink layer necessary
- Routers participate in global delivery
  - Network layer necessary
- Routers don't support reliable delivery
  - Transport layer (and above) **not** supported




42

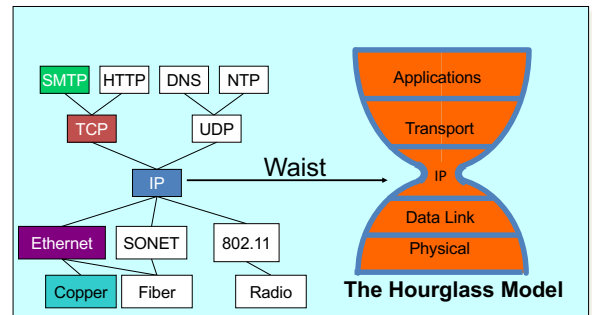


## What Gets Implemented on Switches?

- Switches do what routers do, except they don't participate in global delivery, just local delivery
- They only need to support Physical and Datalink
  - Don't need to support Network layer
- Won't focus on the router/switch distinction
  - Almost all boxes support network layer these days
  - Routers have switches but switches do not have routers



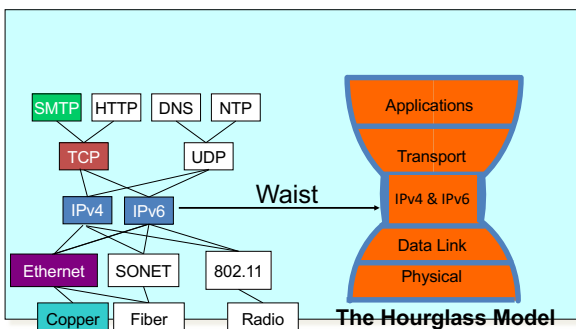
## The Internet Hourglass



There is just **one** network-layer protocol, **IP**.  
The “narrow waist” facilitates **interoperability**.

44

## The middle-age Internet Hourglass



There is just ~~one~~ <sup>TWO</sup> network-layer protocol, **IPv4 + v6**  
The “narrow waist” facilitates **interoperability(???)**.

## Alternative to Standardization?

- Have one implementation used by everyone
- Open-source projects
  - Which has had more impact, Linux or POSIX?
- Or just sole-sourced implementation
  - Skype, Signal, many P2P implementations, etc.

46