Pearls of Computer Science

Computer networks

Pieter-Tjerk de Boer

What's the Internet: "nuts and bolts" view

- millions of connected computing devices:
  - *hosts* = *end systems*
  - running *network apps*

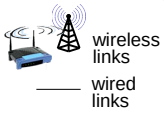

- *communication links*
  - fiber, copper, radio, satellite
  - transmission rate: *bandwidth*

- *Packet switches:* forward packets (chunks of data)
  - *routers* and *switches*

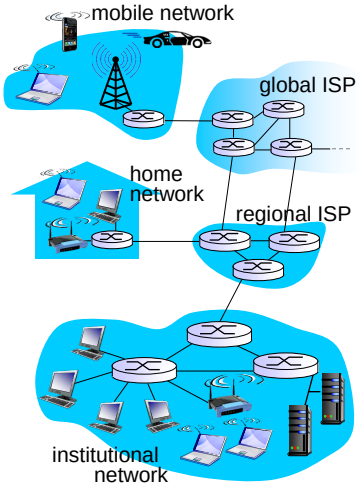

What's the Internet: "nuts and bolts" view

- *Internet:* "network of networks"
  - Interconnected ISPs
- *protocols* control sending, receiving of msgs
  - e.g., TCP, IP, HTTP, Skype, 802.11
- *Internet standards*
  - RFC: Request for comments
  - IETF: Internet Engineering Task Force

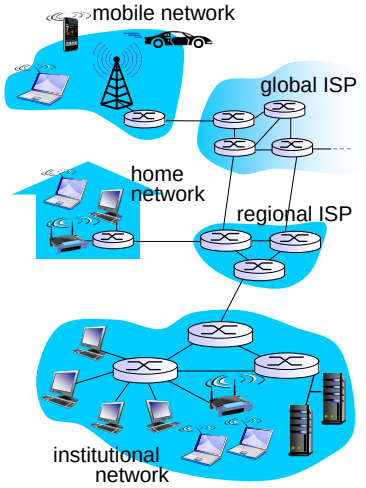

Internet structure: network of networks

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

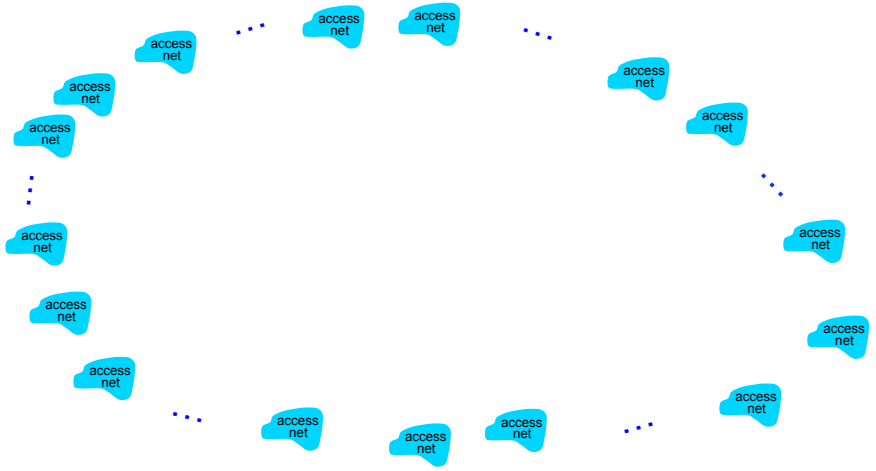

Internet structure: network of networks

Option: connect each access ISP to every other access ISP?

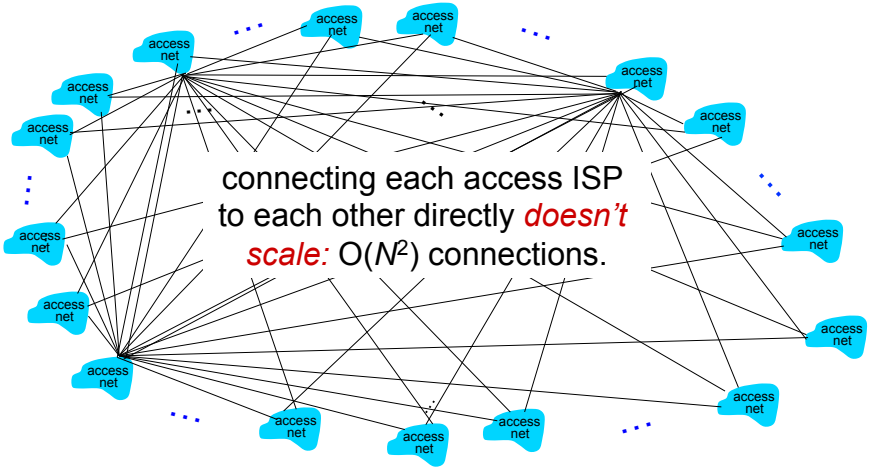

Internet structure: network of networks

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

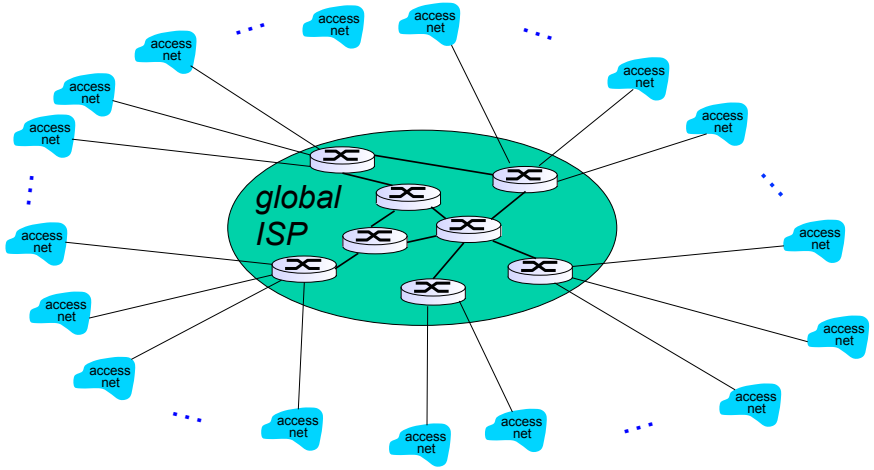

Internet structure: network of networks

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

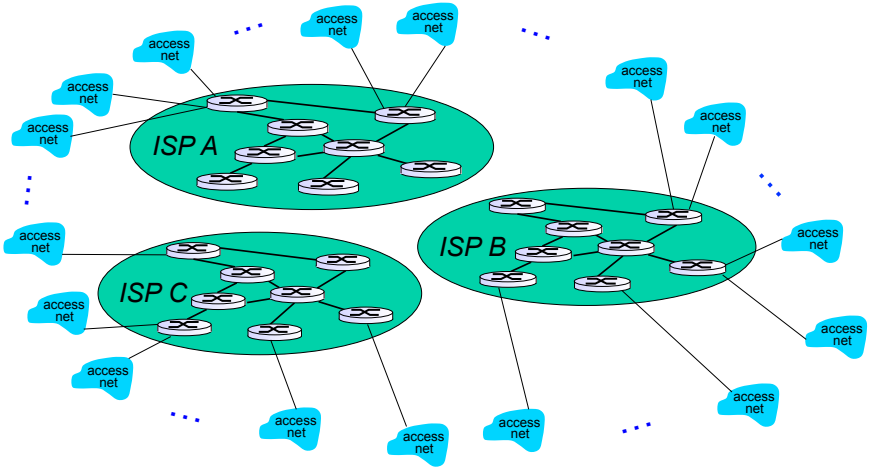

Internet structure: network of networks

But if one global ISP is viable business, there will be competitors …. which must be interconnected

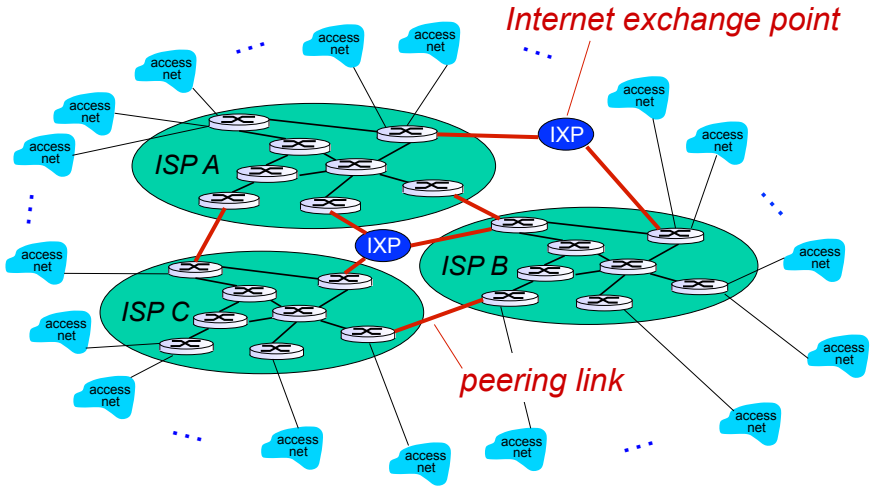

Internet structure: network of networks

… and regional networks may arise to connect access nets to ISPS

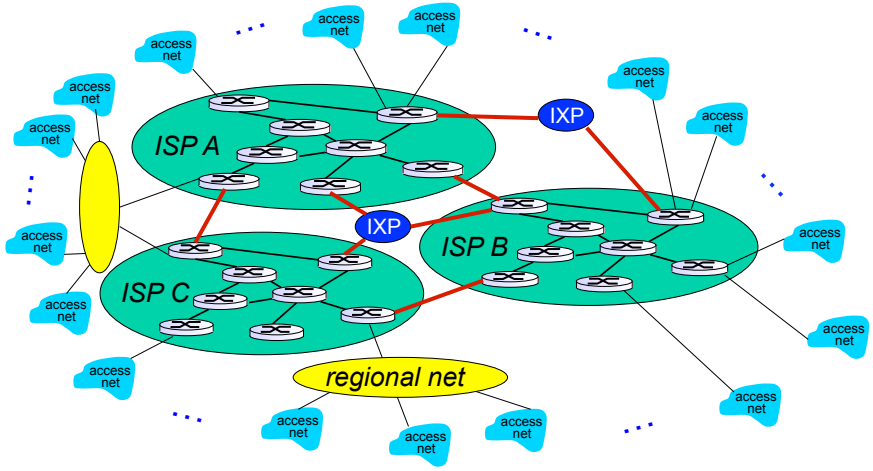

Internet structure: network of networks

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

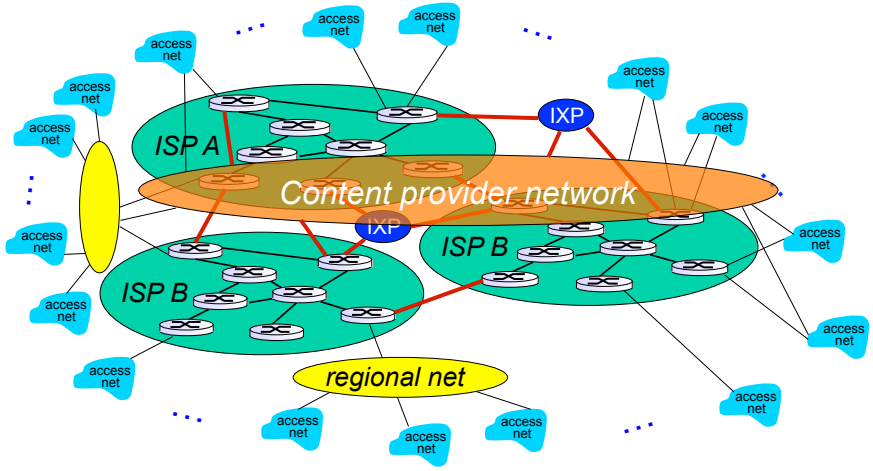

Internet structure: network of networks

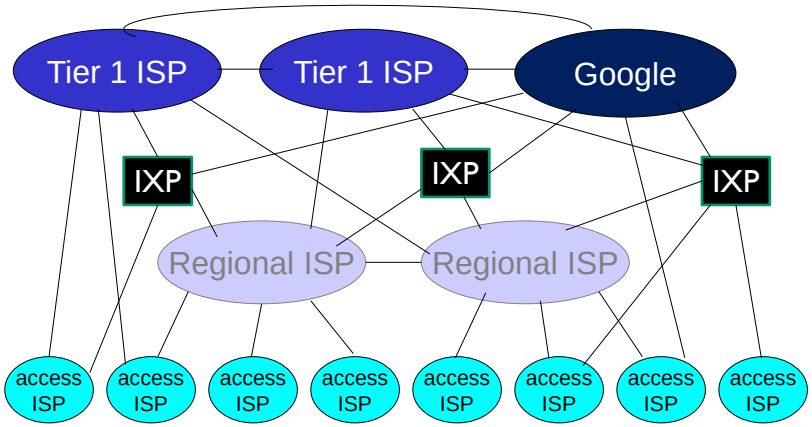

- at center: small # of well-connected large networks
  - "tier-1" commercial ISPs (e.g., Level 3, Sprint, AT&T, NTT), national & international coverage
  - content provider network (e.g, Google): private network that it data centers to Internet, often bypassing tier-1, regional ISPs

Introduction 1-40

Packets vs. circuits

Two methods for organizing communication in a network:

- circuit-switching: for each connection between two endpoints a *part* of the linkspeed is *reserved*.
  Very suitable if the users need a constant number of bits per second; e.g. telephony.
- packet-switching: information flows through the network as packets, which *one after the other* are sent at the *full* link speed.
  Very suitable if the users' needs are very variable; e.g., internet traffic

What is a packet?

- A packet is a block of bits, which are transmitted over a communication link to transport data from one place to another.
- Length typically is variable, but upper-bounded.
- A packet typically contains *data* and a *header*.
- The header contains control information, e.g. about where the data should go.

Layering

- Networked systems often use a *layered* design.
- Each layer *uses* the service from the layer *below* it, to *provide* a better / more useful / more complete service to the layer *above* it.
- Each layer may introduce its own header.
- "One layer's header is (part of) another layer's data."

Internet protocol layers

Layer	Service provided	Internet example
Application	User applications such as web, e-mail, etc.	HTTP, SMTP, etc.
Transport	Reliable delivery of byte stream (Unreliable delivery of packets)	TCP (UDP)
Network	Unreliable delivery of packets troughout a network	IP
Link	Unreliable delivery of packets to neighbour node	Ethernet, WiFi, etc.
Physical	Unreliable delivery of bits to neighbour node	fiber/copper cable, radio channel, etc.

"Hourglass" model

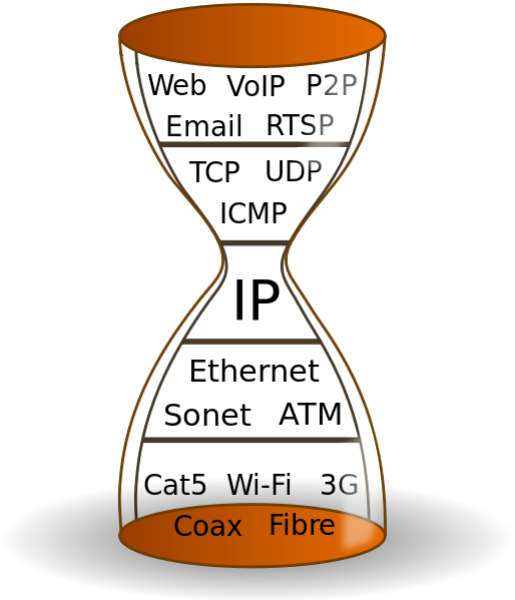

https://commons.wikimedia.org/wiki/File:Internet-hourglass.svg

Encapsulation

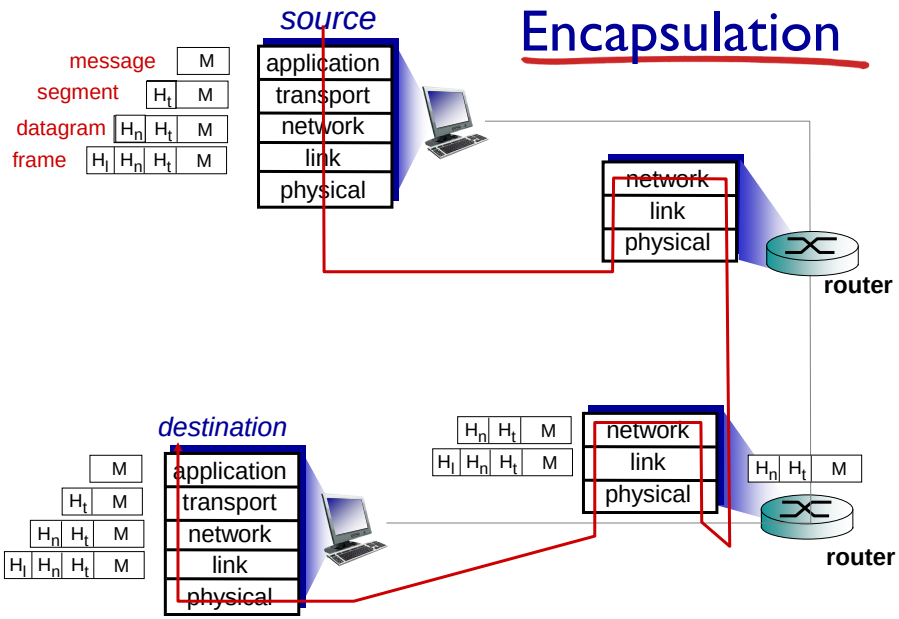

What's a protocol?

a human protocol and a computer network protocol:

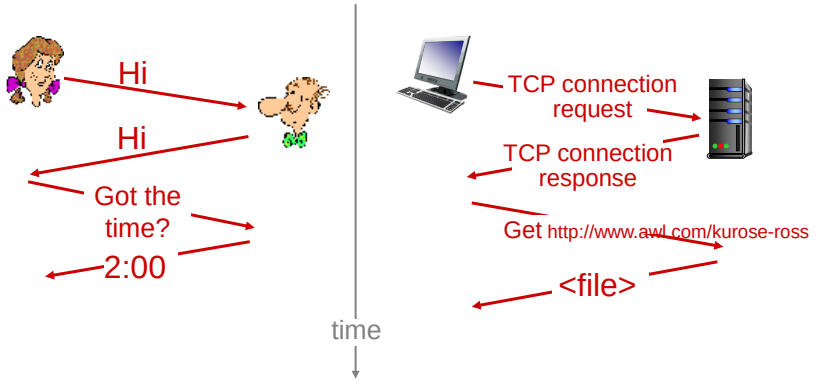

Example protocol for web browsing: HTTP

Request:

```
GET /onderwijs/ewi/index.html HTTP/1.0
Host: www.utwente.nl
User-agent: Mozilla/4.0
```

Reply:

```
HTTP/1.0 200 OK
Content-Length: 12345
Content-Type: text/html

<html>
<head>
  <title>Faculteit EWI</title>
</head>
<body>
  Welkom bij de <b>faculteit EWI</b>!
  Elektrotechniek, Wiskunde en Informatica spelen ...
  ....
  ....
  <img src="foto.jpg">
</body>
</html>
```


Question

Approximately how many computers are connected to the Internet nowadays?

A. 1000 000
B. 10 000 000
C. 100 000 000
D. 1000 000 000
E. 10 000 000 000
F. 100 000 000 000
G. 1000 000 000 000

Adressing

Computers on the internet (“hosts”) are identified by an *address*: a 32-bit number, for example
10000010 01011001 00000011 11111001

For convenience, this is usually written down in decimal, for example 130.89.3.249.

Addresses are assigned *systematically*; e.g., all addresses starting with 130.89 are on the University of Twente.

The above is about IP version 4; version 6 has 128 bits, but is not used much yet.

For human convenience, many hosts also have a *name*, for example `www.utwente.nl` . Converting names to IP addresses and vice versa is done by the Domain Name System (DNS).

Adresses are assigned systematically

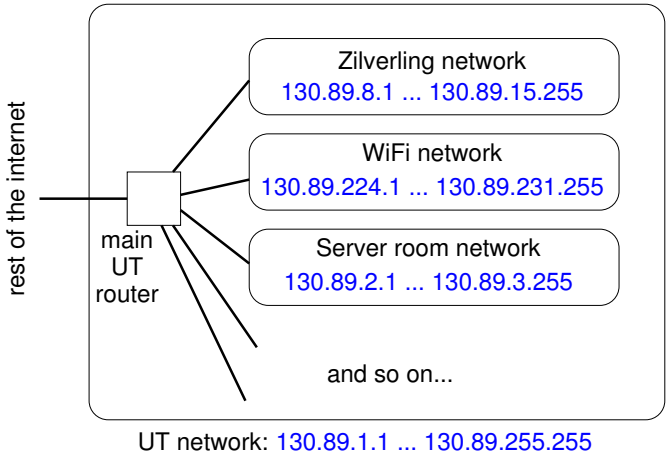

Each sub-network has a subset of the UT’s full address range.

Routers outside the UT only need 1 entry in their forwarding tables for the entire UT; subnetworks are handled within the UT.

Adressing within a host

- Multiple flows of data can exist simultaneously between two hosts.
  ⇒ also at the transport layer a form of addressing is needed.
- Transport layer header contains port number, of 16 bits.
- A TCP connection is uniquely identified by the IP addresses on both sides and the port numbers on both sides.
- Some port numbers are associated with specific applications. Other port numbers are used temporarily.

Dealing with packet loss

Packets can be lost on their way.
Therefore:

- Data packets contain a sequence number.
- Separate packets in the reverse direction confirm correct reception.
- Data packets which are lost, are retransmitted.

Dealing with packet loss

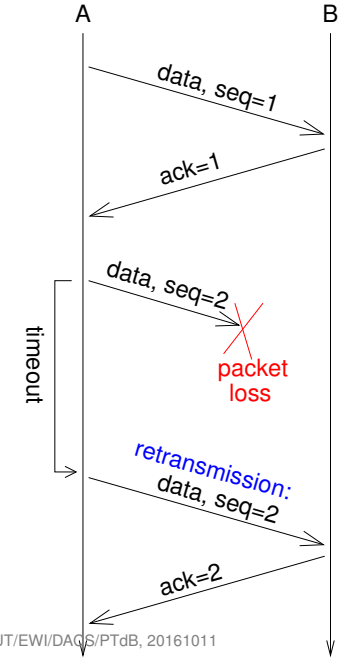

Many variations possible:

- seq.nr. count packets or bytes?
- ack.nr. is last received or next expected?
- can have multiple packets outstanding?
- and if so, retransmit all or one on timeout?
- sender has timeout, or receiver asks sender to repeat?

Needs a protocol!
In the pearl assignment, you'll explore how TCP does this.

How do loss and delay occur?

packets *queue* in router buffers

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

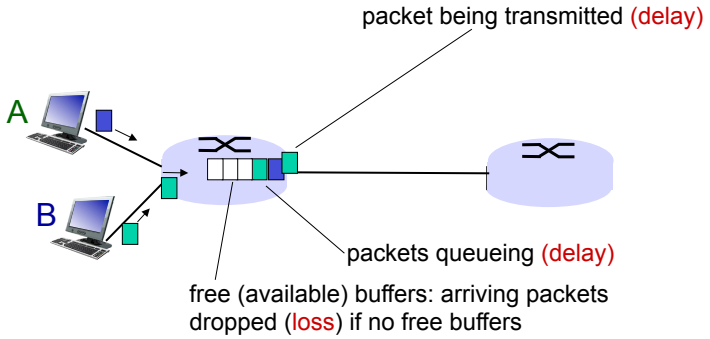

Four sources of packet delay

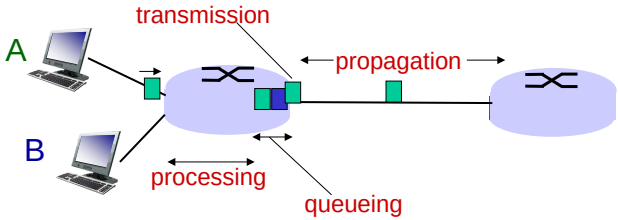

$$d_{total} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{proc}: processing
- check bit errors
- determine output link
- typically < msec

d_{queue}: queueing delay
- time waiting at output link for transmission
- depends on congestion level of router

Four sources of packet delay

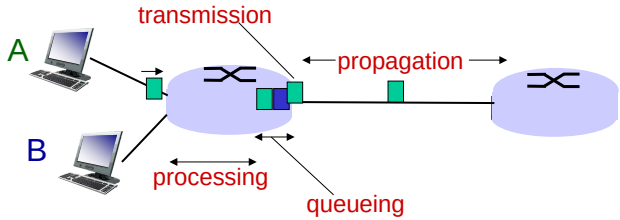

$$d_{total} = d_{proc} + d_{queue} + d_{trans} + d_{prop}$$

d_{trans}: transmission delay:
- L: packet length (bits)
- R: link *bandwidth (bps)*
- $d_{trans} = L/R$

d_{prop}: propagation delay:
- d: length of physical link
- s: propagation speed in medium (~2×10^8 m/sec)
- $d_{prop} = d/s$

d_{trans} and d_{prop} *very* different

* Check out the Java applet for an interactive animation on trans vs. prop delay

Delays for multiple packets

- Be very careful when considering the delay of multiple packets.
- Some delays add up: e.g., next packet can only be transmitted after previous packet's transmission delay has ended.
- Others don't: e.g., several packets can be propagating on the line simultaneously.
  (compare to cars on a road)

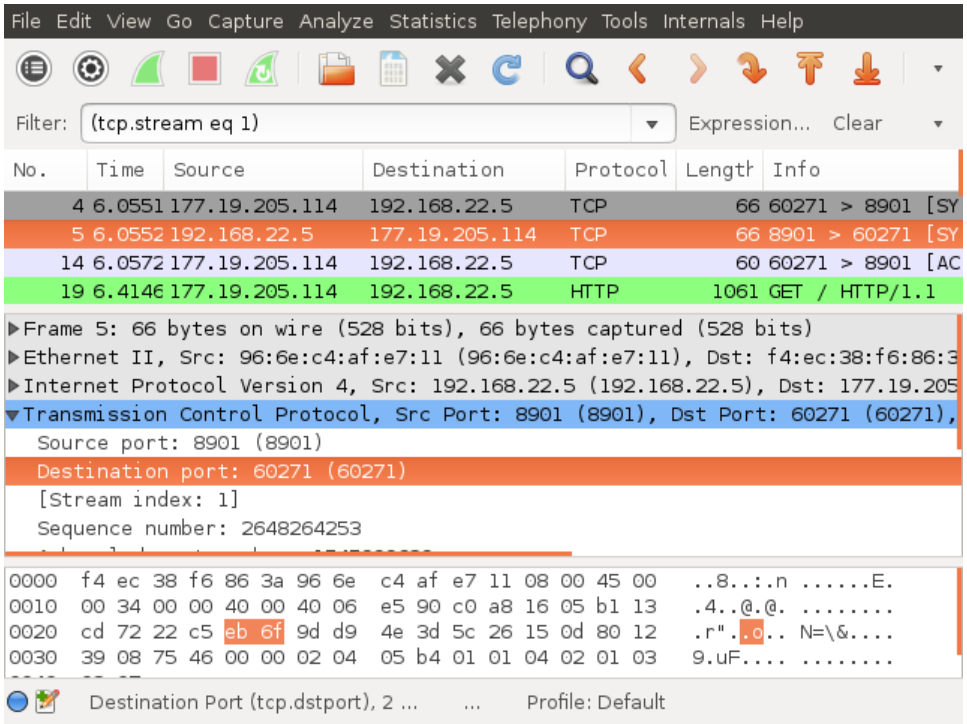

Pearl assignments

- Work with Wireshark: learn more about how the Internet and its protocols work by observing them in action.
- Use traceroute to map paths through the network.

DE VOLKSKRANT
DINSDAG 31 MAART 2009

Dat het Trojaans paard in China staat is niet te bewijzen

Achtergrond

- Ingenieus virus nestelde zich in computers van dalai lama.
- Hackers maakten een fout.

Van onze verslaggever
Bard van de Weijer

AMSTERDAM Op 25 juli 2008 ontvangt een medewerker van het hoofdkantoor van de dalai lama in het Indiase Dharamsala een e-mail die ogenschijnlijk afkomstig is van de bevriende stichting Free Tibet. Aan het bericht zit een document gehecht met informatie voor Tibetanen in ballingschap. Wanneer de medewerker het document aanklikt, opent Microsoft Word en wordt de tekst getoond.

Wat de medewerker niet weet, is dat het bericht niet afkomstig is van Free Tibet, maar van een hacker. De medewerker merkt even-

paard) is zo slim geschreven dat de meeste antivirusprogramma's het niet weten te onderscheppen.

Na verloop van tijd stuiten medewerkers van de dalai lama op vreemde gebeurtenissen. Een buitenlandse diplomaat die per e-mail is uitgenodigd voor een bezoek, krijgt kort daarna een telefoontje van een Chinese functionaris, die hem afraadt af te reizen.

De organisatie vraagt daarop onderzoekers van onder meer het Canadese Munk Center for International Studies te kijken naar wat er aan de hand is. Uit hun onderzoek blijkt dat onder meer kantoren van de dalai lama in Londen, New York, Brussel en India zijn gehackt. Daarnaast zijn computers van bevriende Tibet-organisaties gekraakt. Maar ook het ministerie van Buitenlandse Zaken van Iran, de ambassade van Roemenië en de ambassade van Zuid-Korea in China blijken geïnfecteerd.

besmetting zit. Om dat te achterhalen, installeren de onderzoekers op getroffen computers het contraspionageprogramma Wireshark. Het blijkt dat gevoelige documenten worden verzonden naar vier centrale internetcomputers, *control servers*. Met deze servers, die voornamelijk in China staan, kunnen hackers de controle overnemen over alle computers in het besmette netwerk.

De indringers hebben echter een fout gemaakt, waardoor de onderzoekers zich toegang kunnen verschaffen tot de *control servers*. Zij kunnen nu op hun beurt, zonder dat de hackers het weten, volgen wat er gebeurt.

Veel van de waarnemingen wijzen erop dat de Chinese overheid achter de kraak zit. Maar, stellen de onderzoekers, directe betrokkenheid van Peking is niet te bewijzen. Op internet is het voor kwaadwilligen nog altijd vrij eenvoudig hun