Introduction to Internet QoS and Network Performance

François Fluckiger
CERN, Geneva

Part 1

Internet Basics, TCP behaviour, QoS

Part 2

Transporting Multimedia over the Internet

Epilog

Who has a good or fair understanding of the difference between …

FTP *and* **HTTP**

HTTP *and* **TCP**

TCP *and* **IP**

TCP *and* **UDP**

UDP *and* **RTP**

Internet: Back to Basics

- **Formed of switches called routers**
- **Routers interconnected by links**
- **Topology usually meshed**
- **Hosts chop data stream into blocks called packets**
- **Routers switch individual packets**

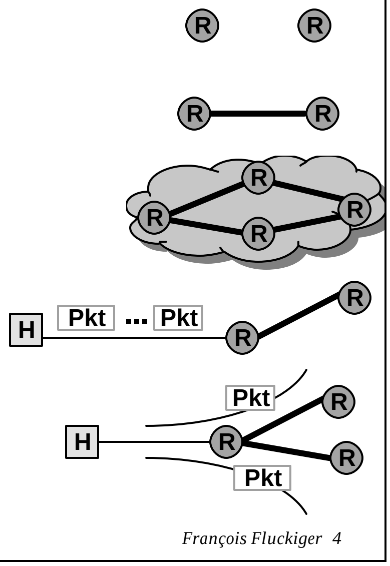

IP

- **Layer 3 (Network) protocol**
- **Specifies**
  - format of packets (size, header, …)
  - mechanism for routing
  - resulting service
- **The highest level protocol understood by routers**

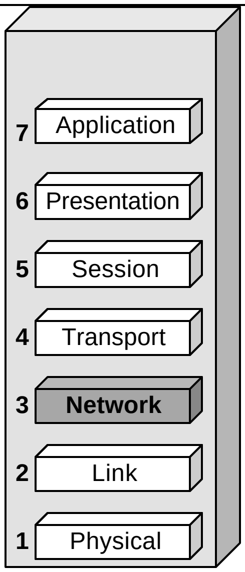

7 Application

6 Presentation

5 Session

4 Transport

3 **Network**

2 Link

1 Physical

Protocols understood by routers

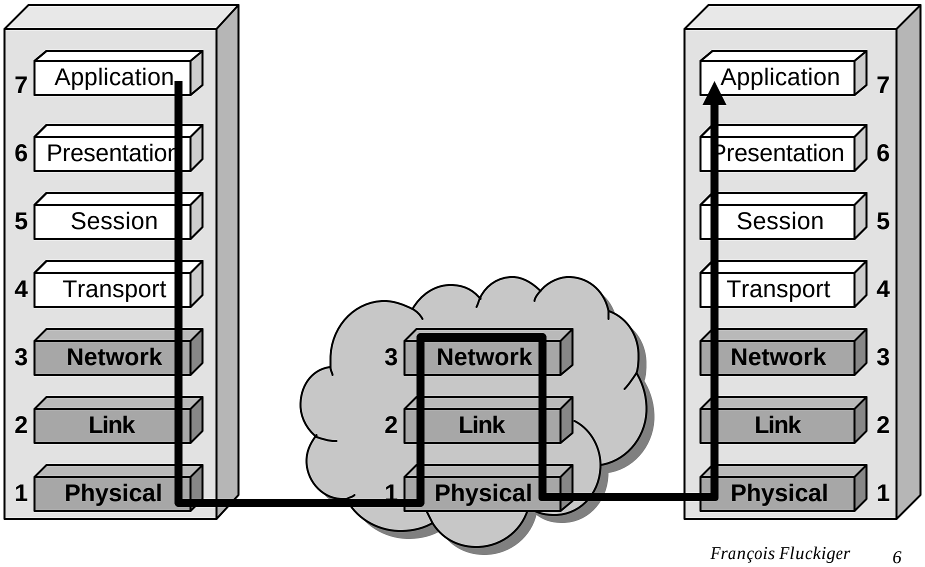

7 Application 7

6 Presentation 6

5 Session 5

4 Transport 4

3 **Network** 3 **Network** 3 **Network** 3

2 **Link** 2 **Link** 2 **Link** 2

1 **Physical** 1 **Physical** 1 **Physical** 1

Layering principle (emission)

- Each protocol layer **N** adds a **Header** to the data unit received from layer **N+1** (1)
- Header contains control information; e.g. :
  - Numbering of the data unit
  - Coding of the destination
  - Codes for error detection
  - Priority of the data unit

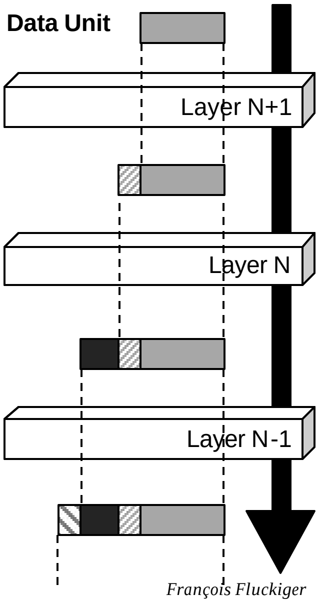

(1) and segment the data unit if necessary

Packets

- Each **data unit** generated at a given layer has a specific name
- Data Unit at layer 3 (IP) is called a **Packet**

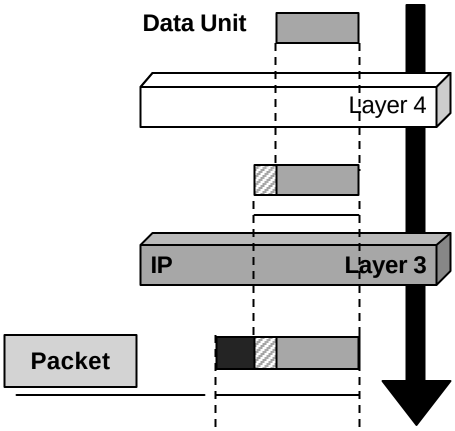

Main IP features

- **IP is a connectionless (CL) protocol**
  - all packets independently routed
  - packets carry full destination address
  - **all packets have same priority**

- **Opposite = connection-oriented (CO)**
  - no information sent before a hard connection is set up

CO and CL networks: Examples

Connection-oriented	**Connectionless**
■ Telephone	■ Post office ■ Road Network
■ ATM ■ Frame Relay ■ *SNA* ■ *X.25*	■ All regular LANS ■ Ethernet ■ *FDDI* ■ *Token Ring* ■ **Internet IP**

CO vs. CL

+

CO	CL
■ Traffic more predictable	■ No call set-up delay before sending a packet
■ Easier for network to reserve resources	■ Routing possibly more dynamic
■ QoS guarantees easier to provide	■ Resilience

TCP

- ■ **Layer 4 (Transport) protocol**
- ■ **Specifies**
  - ■ Format of segment
  - ■ Mechanisms for flow control, error detection, error recovery
- ■ **The lowest protocol used by hosts**

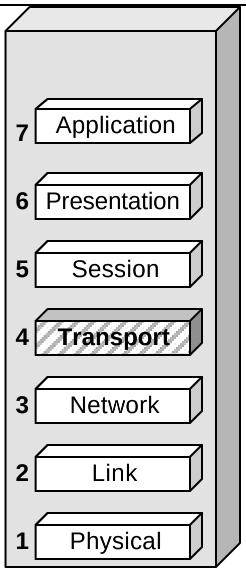

Protocols understood by hosts only

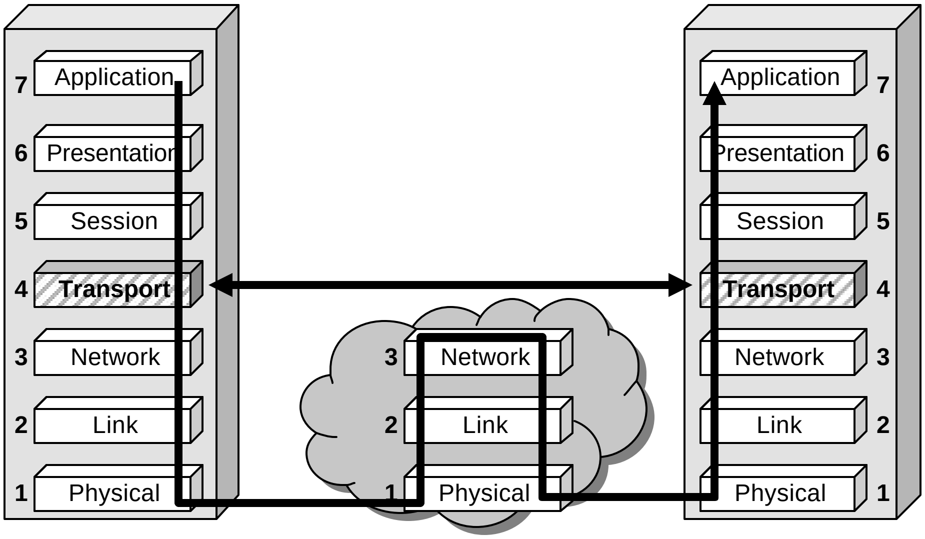

TCP Segments

- Each **data unit** generated at a given layer has a specific name
- Data Unit at layer 3 (IP) is called a **Packet**
- Data Unit at layer 4 (TCP) is called a **Segment**

  (sometimes also called **block**)

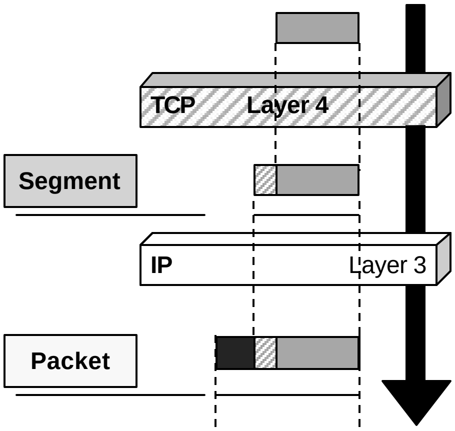

What does TCP Provides

- **Error services**
  - Detection of corrupted data
  - Detection of loss, duplicated, out of sequence packets
  - Correction of errors
- **Flow control between receiver/server**
- **Mechanisms to limit network congestions**

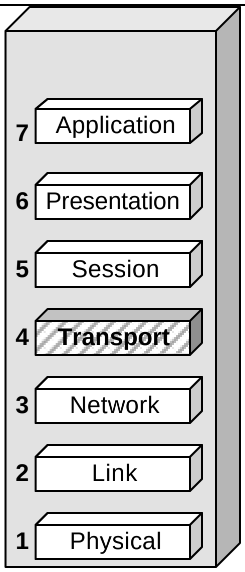

TCP Mechanisms for Errors

- **Segments**
  - Theoretical maximum= 65535 bytes
- **Error detection**
  - Segments carry a **sequence number**
  - Sequence number = the **order of the first octet of segment** in the data stream
  - Receiver can detect out-of-sequence packets

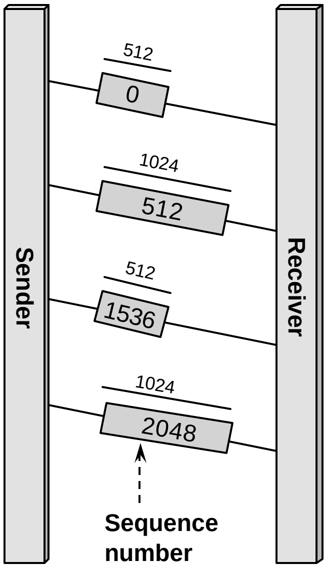

TCP Mechanisms for Errors

- **Error notification**
  - Positive ACK by receiver
  - e.g. "ACK 512 means":

    " *I am now ready to receive octet #512 and beyond, because I correctly received all octets up to 511*"
- **Problem with this**
  - Sender needs to wait for ACK of segment 1 before sending segment 2

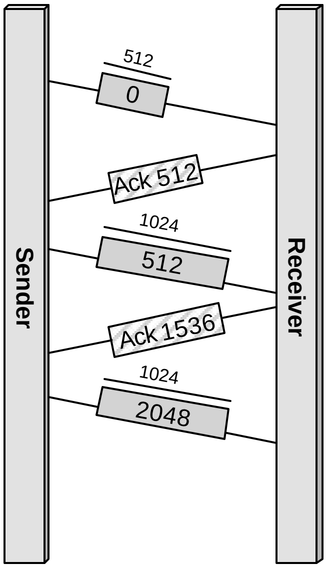

TCP Windowing

- **Windowing**
  - Mechanism to anticipate on the ACK
  - Sender manages a byte count which gives the limit of the highest octet that can be sent without being acknowledged
  - This is called the **Congestion Window** (**CWND**)

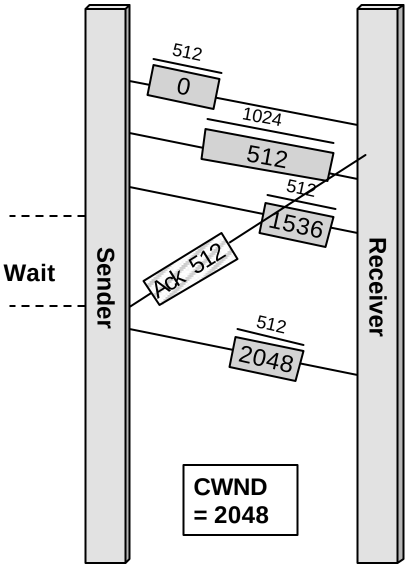

TCP Sliding Window

Initial window

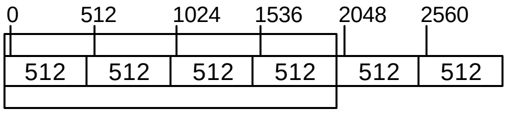

Initial window after 1st ACK

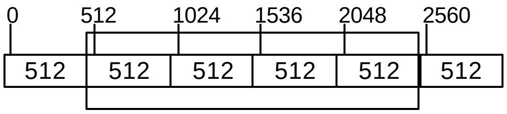

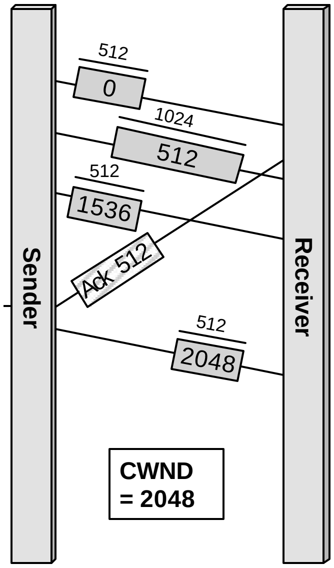

TCP mechanism for Flow control

- **Dynamic Window size**
  - Window size CWND not fixed
  - Sender may reduce if detecting network congestion
  - Receiver indicates value of window in ACK
  - Value of Window = remaining incoming **buffer space** in receiver

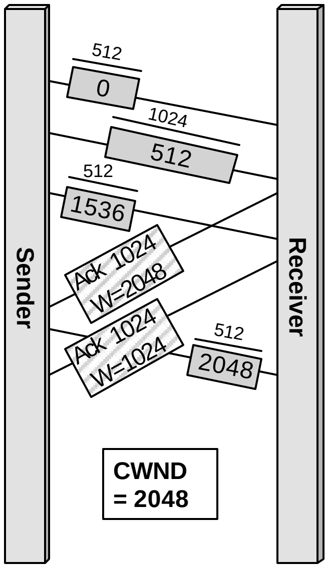

TCP mechanism for Flow control

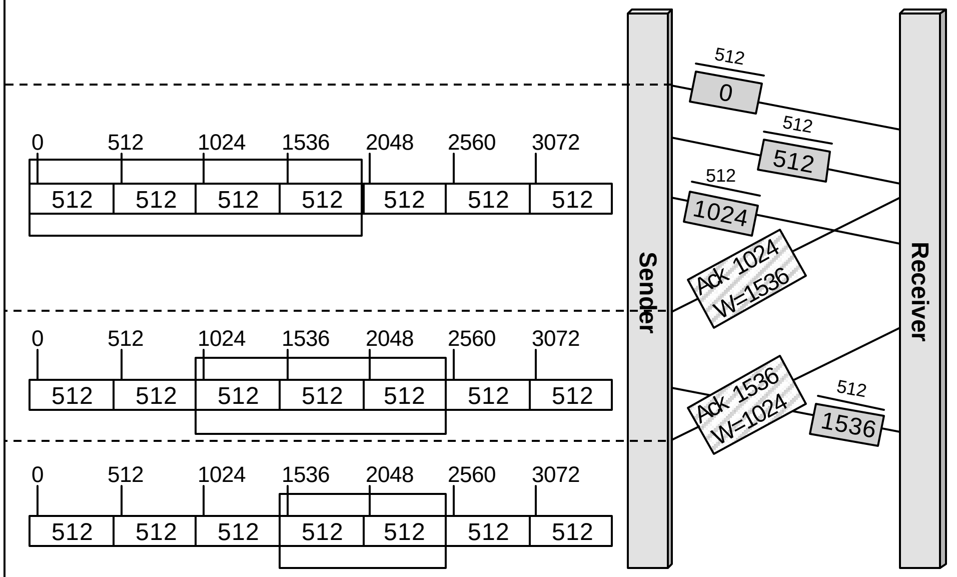

TCP buffer size

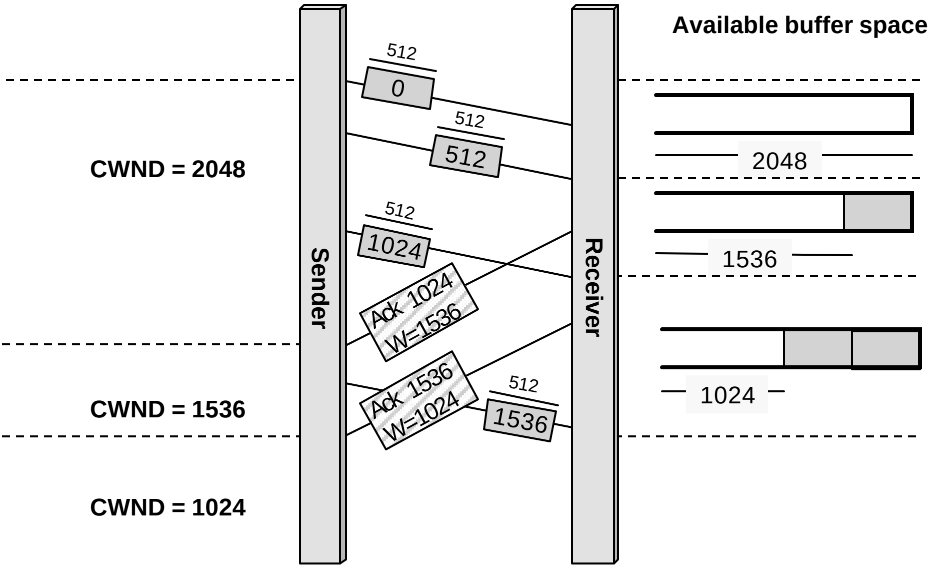

TCP mechanism error correction

- **TCP retransmits non-acknowledged segment**
- Question: ***when to retransmit***?
- **Timer: TCP …**
  - starts a timer on any transmission
  - waits for the ACK
  - retransmits segment in no ACK when timer expires

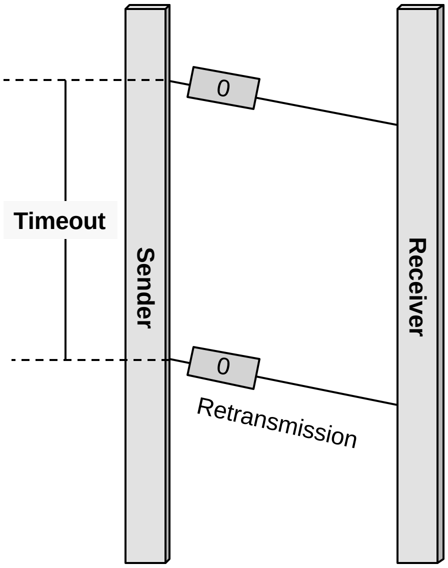

TCP Timeout

- **Question: *which value for the Timeout*?**
- **TCP** measures permanently the **Round Trip Time (RTT)**
  - RTT = Average Time between segment emission and ACK reception
- **Timeout = b x RTT**
  - Choice of b delicate (simple choice: b = 2)

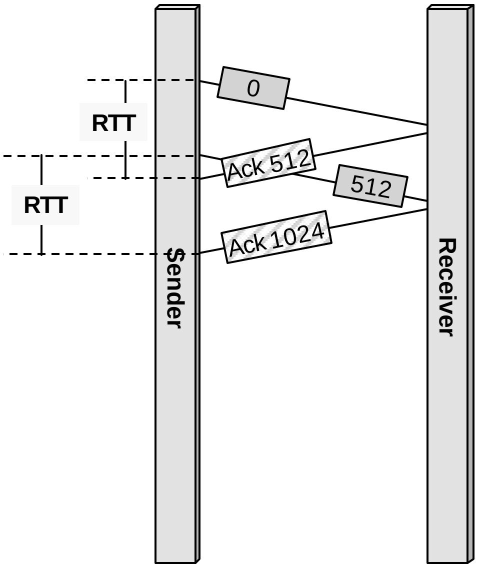

Optimal buffer size

- **Sender must keep all sent segments until acknowledged**
- **Question:**

  What is the optimal buffer size to keep all segments?

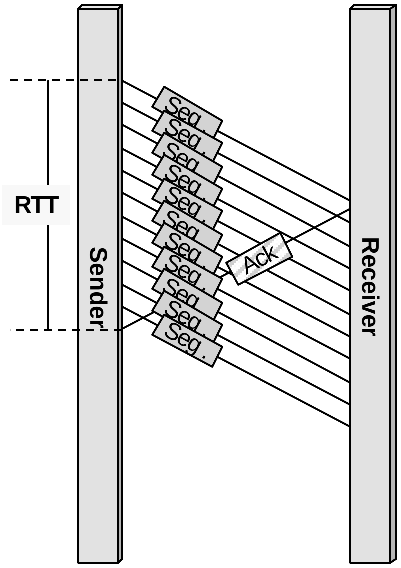

Optimal buffer size

- If bottleneck of the Bandwidth between Sender and Receiver is **B** byte/second
  - Sender can send up to **B** bytes / second over **RTT** seconds
  - Maximum number of bytes to store = **B x RTT**
- Called the **Bandwidth*Delay Product** (**BDP**)

  BDP = B x RTT

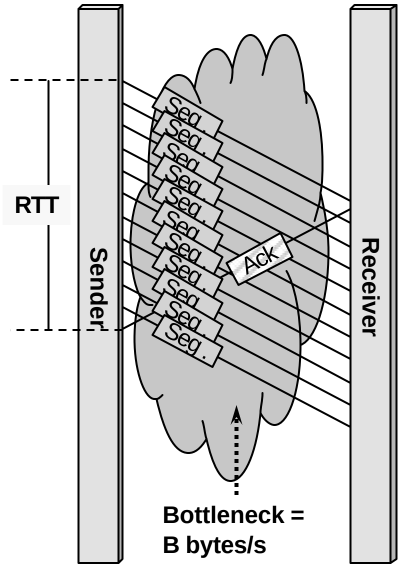

Optimal buffer size

- **Optimal buffer size** (1)

  =

  Bandwidth*Delay Product

- **BDP = B x RTT**

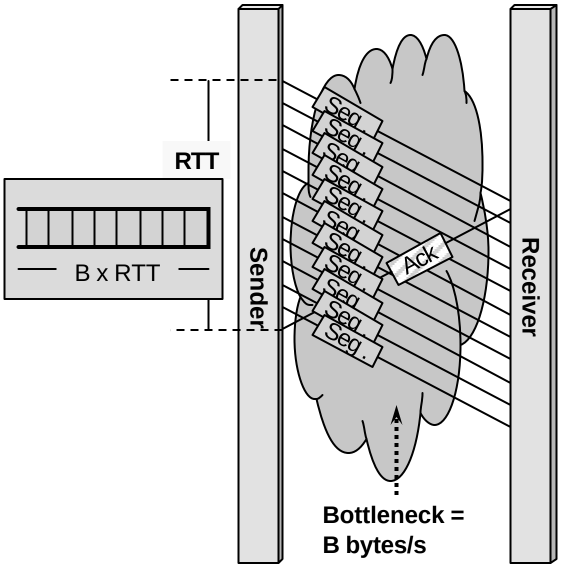

- (1): sender and receiver

Avoiding congestions

- When TCP starts sending, ignores the network bottleneck (often at L|AN –WAN interface)
- Can't send a full speed (would flood the network).
- Start with
  - maximum segment size, but
  - with minimum congestion window (CWND) = 1 segment
- On every ACK, increases CWND by 1

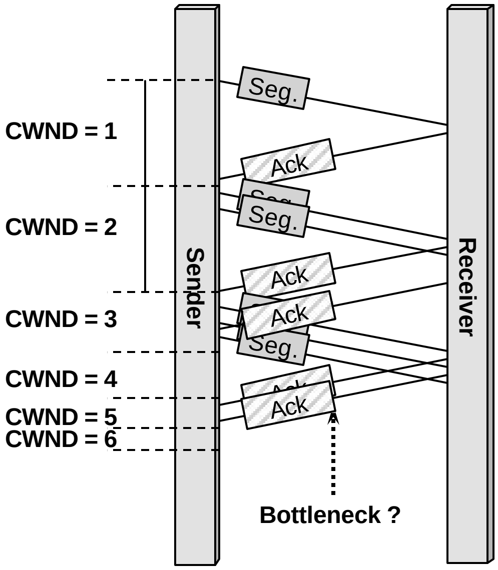

TCP congestion avoidance

- **Called slow start**
  - Rate **doubles** every round-trip time
- **If packet loss, sender halves window**
  - Then window increases **linearly**
- **In case of end-to-end collapse** (sender time out exhausted), **sender resumes with slow start**

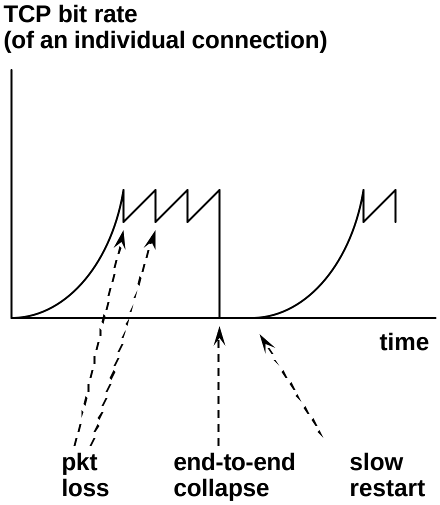

Congestion Notification in Internet

- **Current Internet, no explicit congestion notification from network to source**
- **Notification technique:** ***Network drops packets***
- **TCP will react and slow down**
- **TCP is said to be** ***congestion-indication responsive***

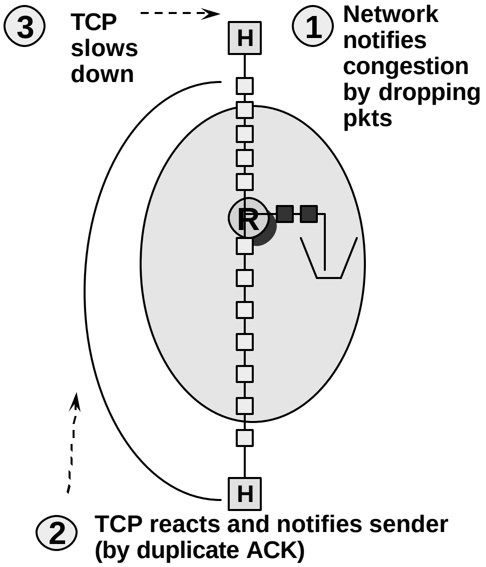
3
TCP slows down
1
Network notifies congestion by dropping pkts
H
R
H
2
TCP reacts and notifies sender (by duplicate ACK)

Internet Base IP service

- **Initial Internet**
  - single class of service:“***best-effort*** **service**”
  - packet forwarding completely **egalitarian**
  - **No service guarantee**
- **How to better guarantee end-to-end throughput, delays?**
  - ***"How to have packet more equal than others?"***

On service discrimination ...

- **Objective**
  - Give **better** service to some traffic
- **Consequence**
  - **…** at the expense of giving **worse** service to the rest *(hopefully in times of congestion only)*

Improving Internet QoS

Two proposed techniques

- ***Integrated Services***
- ***Differentiated Services***

Integrated Services principles

- ***Resource reservation*** **is necessary**
- **Reservations on a per-flow basis**
- **Routers have to maintain flow-specific *states***
- ***Protocol: RSVP***
- ***Problem: does not scale well***

Differentiated Services : principles

- **Rationale**
  - With RSVP, knowing pkt priority needs heavy classification process
  - Classification may be faster if packets are explicitly "***marked***" (1) with a priority

(1) "Marking" pkts also called "***coloring***"

Part 2

Multimedia over the Internet and QoS

Real-time media transmission (A / V)

- Assume now the source of packets in an
  - analog digital converter
  - connected to a microphone and a movie camera
- Called **Streaming** audio/video

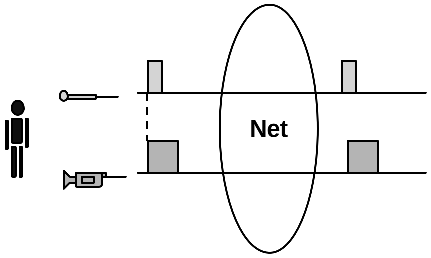

Audio and video packets

TCP behaviour

- **Slow start**
- **Sender aware of packets dropped**
- **Sender decreases bit rate when packet dropped**

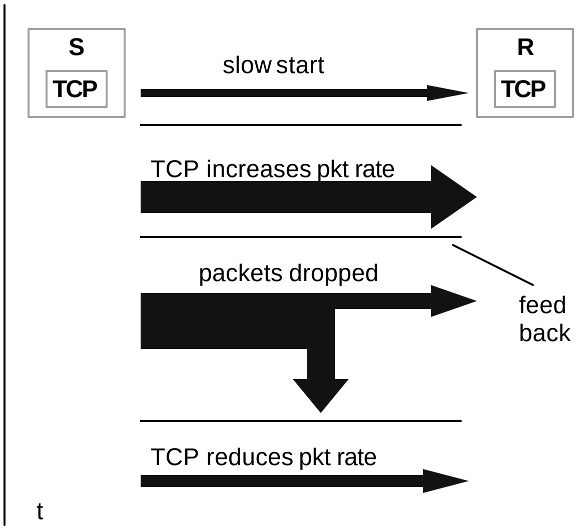

UDP behaviour

- **UDP sends blindly to a receiver**
- **No feedback from the receiver**
- **Sender unaware whether packets are dropped/lost**

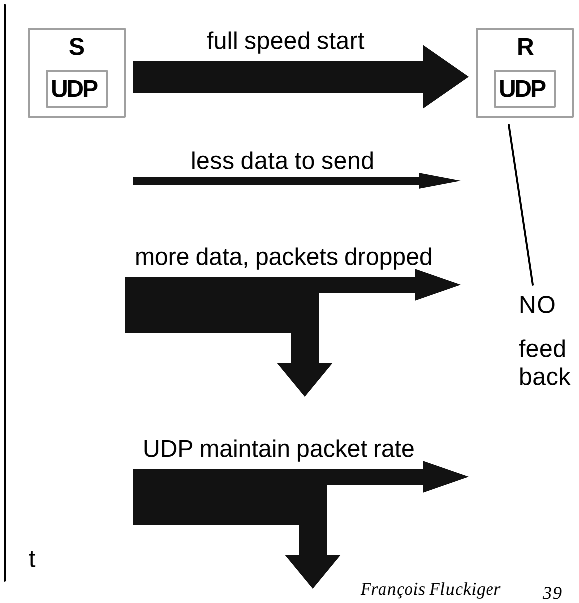

Protocols for real-time audio and video

- **Audio/video applications cannot operate over TCP**
  - slow start, congestion control, ...
- **They use UDP**
  - But UDP has no feedback, ...
- **Most applications use RTP (Real-Time Transport Protocol)**
  - packet loss detection
  - Some form of feedback

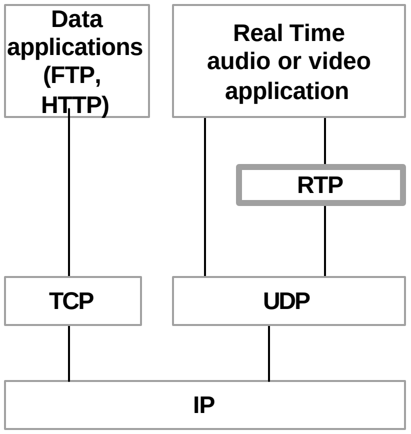
Data applications (FTP, HTTP)
Real Time audio or video application
RTP
TCP
UDP
IP

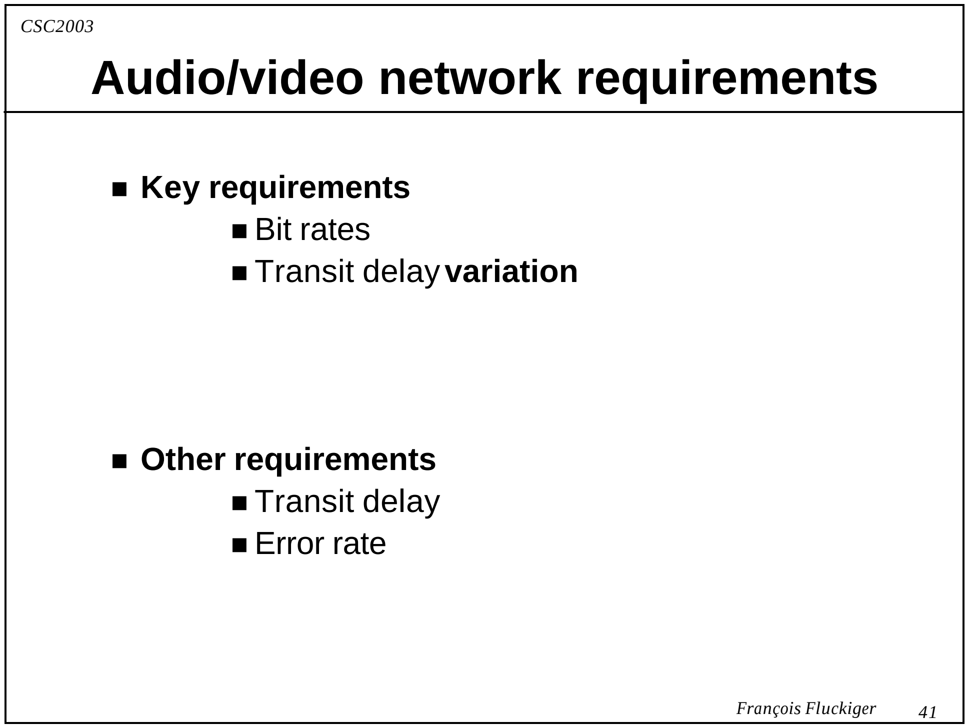

Audio/video network requirements

- **Key requirements**
  - Bit rates
  - Transit delay **variation**

- **Other requirements**
  - Transit delay
  - Error rate

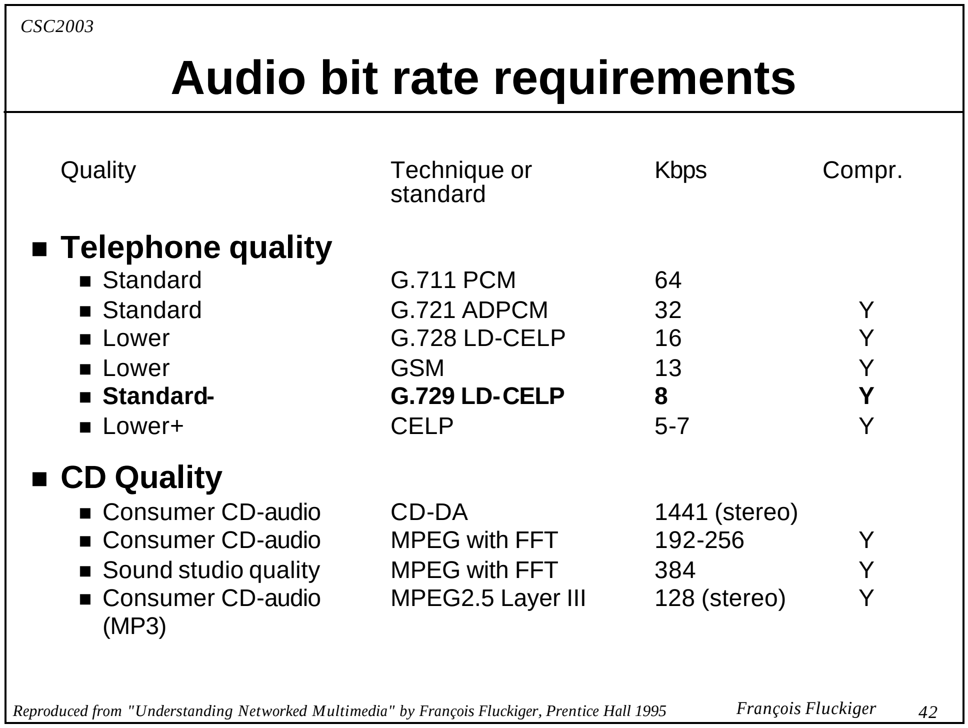

Audio bit rate requirements

Quality	Technique or standard	Kbps	Compr.
■ **Telephone quality**			
■ Standard	G.711 PCM	64	
■ Standard	G.721 ADPCM	32	Y
■ Lower	G.728 LD-CELP	16	Y
■ Lower	GSM	13	Y
■ **Standard-**	**G.729 LD-CELP**	**8**	**Y**
■ Lower+	CELP	5-7	Y
■ **CD Quality**			
■ Consumer CD-audio	CD-DA	1441 (stereo)	
■ Consumer CD-audio	MPEG with FFT	192-256	Y
■ Sound studio quality	MPEG with FFT	384	Y
■ Consumer CD-audio (MP3)	MPEG2.5 Layer III	128 (stereo)	Y

Reproduced from "Understanding Networked Multimedia" by François Fluckiger, Prentice Hall 1995

Which bit rate is actually needed?

- **Network overheads incl.:**
  - RTP header (12 bytes)
  - Transport Protocol header (usually UDP, 8 bytes)
  - IP header (20 bytes)
  - Example:

    raw G.711 **64** Kbps requires from **68** to 80 Kbps

- ***However, speech contains silence***

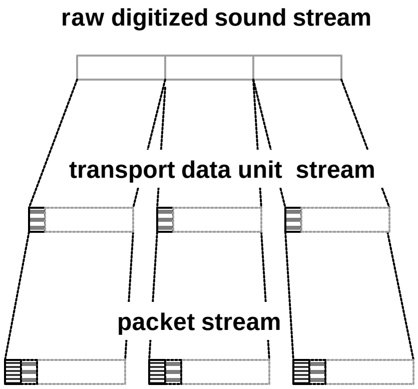

Silences in speech

- **Monologue**
  - typically **15%** silence
- **Bi-party telephone conversation**
  - **20%** silence for overall conversation
  - **60%** silence for each party
  - If silence suppressed, required bit rate is in effect **<40%** of nominal raw bit rate

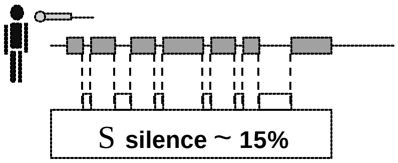

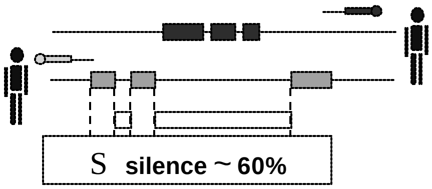

Video bit rate requirements

Quality	Technique or standard	Mbps	Compr.
■ **Video conf. quality**	**H.261**	**0.1**	**Y**
■ **VCR quality**	**MPEG-1**	**1.2.**	**Y**
■ **Broadcast quality**			
■ **Compressed**	**MPEG-2**	**2-4**	**Y**
■ **Compressed**	**MPEG-4**	**2**	**Y**
■ **Studio-quality digital TV**			
■ **Uncompressed**	**ITU-R 601**	**166**	
■ **Compressed**	**MPEG-2**	**3 to 6**	**Y**
■ **HDTV**			
■ **Uncompressed**	**CD-DA**	**2000**	
■ **Compressed**	**MPEG-2**	**25 to 34**	**Y**

Reproduced from "Understanding Networked Multimedia" by François Fluckiger, Prentice Hall 1995

Principle of Compression

- **Remove redundancies from the original signal**

Information rate, bit rate, entropy

- **Information content or *entropy* of a sample:**
  - a function of how different it is from the predicted value

- **Shannon's theory:**
  - any signal which is **totally predictable**: carries **no information**
    - (e.g. a sine wave)
  - **noise** is completely **unpredictable:** high entropy

Real, contentless signals, noise

Real signals

lie somewhere between the two extremes

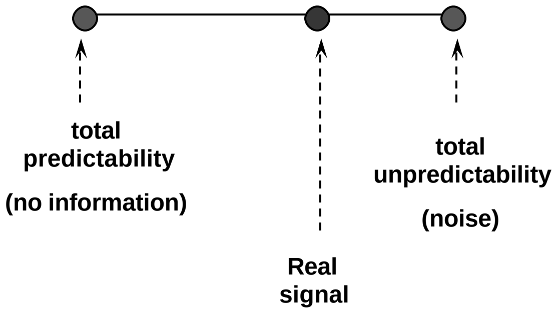

The Effect of compression

Compression removes redundancy ... *but*

The Effect of compression

- **Redundancy essential for resistance to errors**
- **Compressed data more sensitive to errors**

Unresponsive flows

- **Unresponsive flows do not react to congestion indication (pkt loss)**
- **Can create**
  - bandwidth **starvation** inflicted to well-behaved responsive traffic

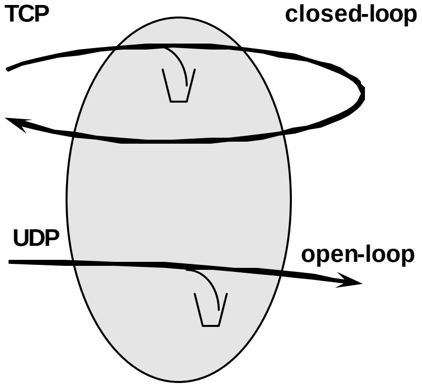

Unfair competition

- **Case 1:**
  - **2 LANs (10 Mbps)**
  - **interconnected with T1 and a pair or routers**
- **Competition between**
  - **3 TCP connections and**
  - **1 UDP connection**

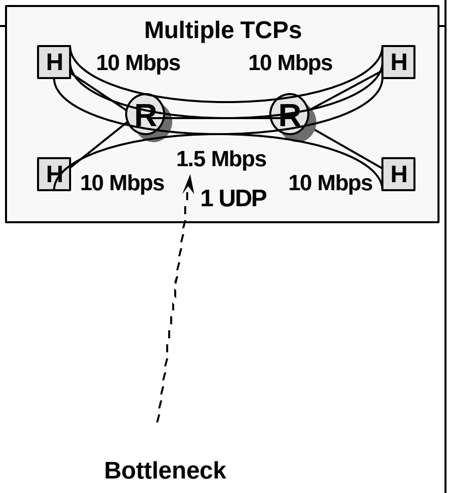

From S.Floyd et al, February 97

Unfair competition

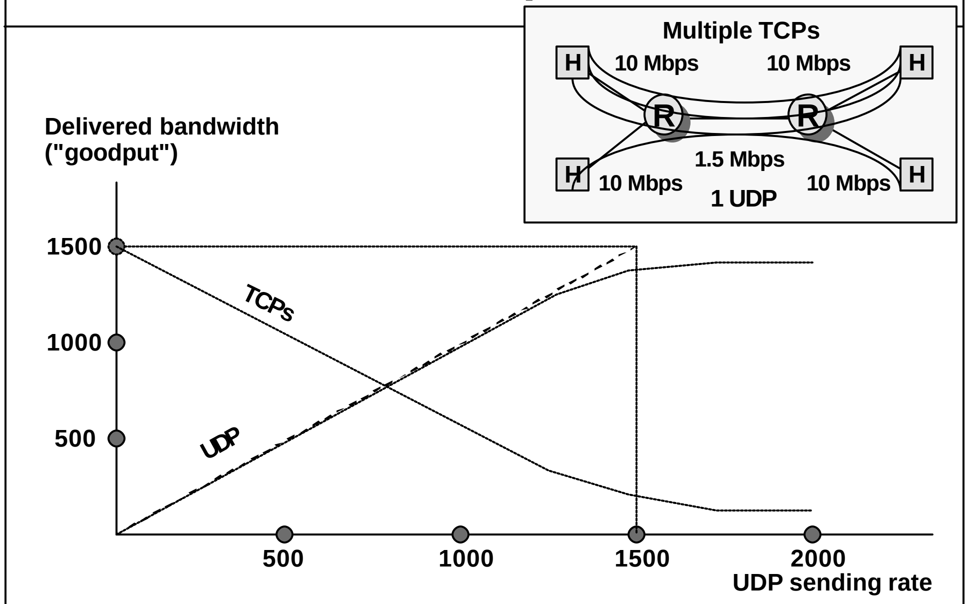

From S.Floyd et al, February 97

Further reading ...

- **Internetworking with TCP/IP, ol 1**

  Douglas E. Commer, Prentice Hall, ISBN 0-130-183806
- **Computer Networks, Ed. 4**

  Andrew Tannenbaum, Prentice Hall, ISBN 0-130-661023
- **Understanding Networked Multimedia**

  Francois Fluckiger, Prentice Hall, ISBN 0-131-90992-4
- **Understanding Media**

  Marshal Mac Luhan, The MIT Press, ISBN 0-262-631159