Internet Transport Layer

TCP Fundamentals, TCP Performance Aspects,
UDP (User Datagram Protocol)

Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - Slow Start and Congestion Avoidance
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- **RFC Collection**

Internet Transport Layer, v4.5   2

Transport Layer Protocols

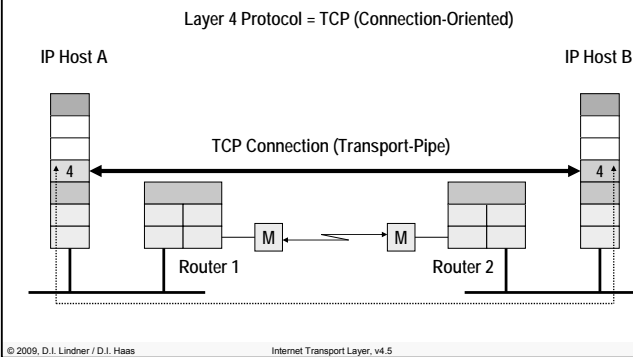
Internet Transport Layer, v4.5   3

TCP (Transmission Control Protocol)

- **TCP is a connection oriented layer 4 protocol (transport layer) and is transmitted inside the IP data field**
- **Provides a secure end-to-end transport of data between computer processes of different end systems**
- **Secure transport means:**
  - Error detection and recovery
  - Maintaining the order of the data without duplication or loss
  - Flow control
- **RFC 793**

Internet Transport Layer, v4.5   4

TCP and OSI Transport Layer 4

Layer 4 Protocol = TCP (Connection-Oriented)

IP Host A

IP Host B

TCP Connection (Transport-Pipe)

4

4

M

M

Router 1

Router 2

Internet Transport Layer, v4.5

5

TCP Protocol Functions 1

- **Data transmission within segments**

- **ARQ protocol with Continuous Repeat Request technique and piggy-backed acknowledgments**

- **Error recovery through sequence numbers (based on octets !), positive & multiple acknowledgements and timeouts for each segment**

- **Flow control with sliding window and dynamically adjusted window size**
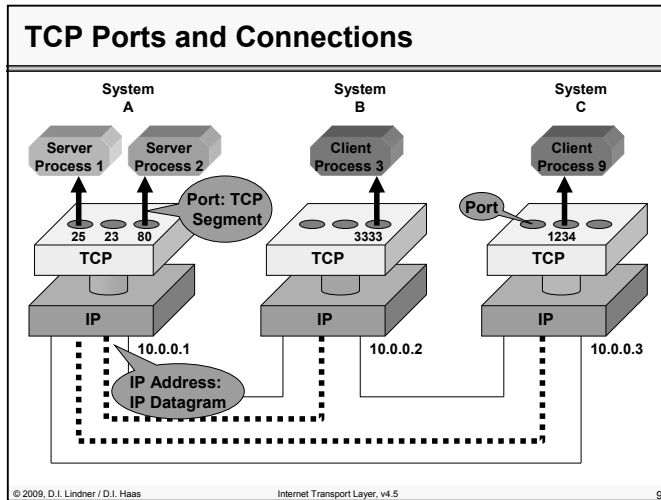
Internet Transport Layer, v4.5

6

TCP Protocol Functions 2

- **In general, segments are encapsulated in single IP packets**
- **Maximum segment size depends on max. packet or frame size (fragmentation is possible)**

- **Call setup with "three way handshake"**

- **Hides the details of the network layer from the higher layers and frees them from the tasks of transmitting data through a specific network**

Internet Transport Layer, v4.5

7

TCP Ports and Connections

- **TCP provides its service to higher layers through ports**
- **Each communicating computer process in an IP host is assigned a port**
  – identified by a port number
- **By usage of ports the TCP software can serve multiple processes (browser, e-mail etc.) simultaneously**
- **The TCP software functions like a multiplexer and demultiplexer for TCP connections**
  – Port 25 on system A:
    • process 1, system A <--------> process 3, system B
  – Port 53 on system A:
    • process 2, system A <--------> proc. 9, system C
  – (see next slide)

Internet Transport Layer, v4.5

8

TCP Ports and Connections
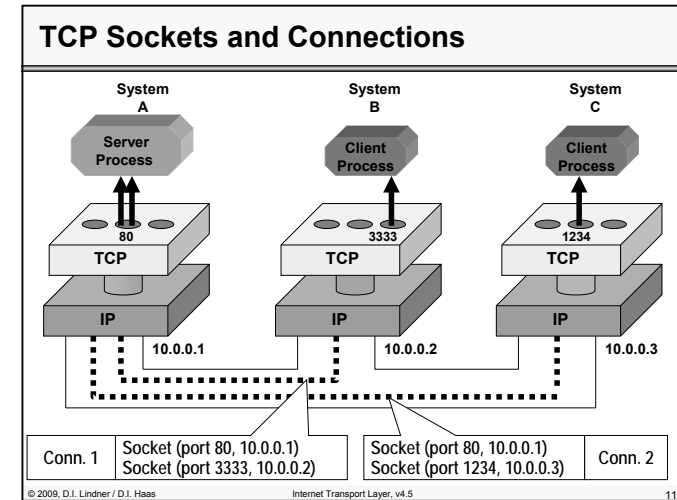
TCP Sockets and Connections

- **In a client-server environment a communicating server-process has to maintain several sessions (and also connections) to different targets at the same time**
- **Therefore, a single port has to multiplex several virtual connections**
  - these connections are distinguished through sockets
- **The combination IP address and port number is called a "socket"**
  - Each socket pair uniquely identifies a connection

TCP Sockets and Connections

TCP Well Known Ports

- **Well known ports**
  - Are reserved for common applications and services (like Telnet, WWW, FTP etc.) and are in the range from 0 to 1023
  - Are controlled by IANA (Internet Assigned Numbers Authority)
- **Registered ports**
  - start at 1024 (e.g. Lotus Notes, Cisco XOT, Oracle, license managers etc.). They are not controlled by the IANA (only listed, see RFC1700)

- **Well-known ports together with the socket concept allow several simultaneous connections (even from a single machine) to a specific server application**
- **Server applications listen on their well-known ports for incoming connections**
- **Port concept und port numbers also used for UDP**

L11 - Internet Transport Layer (TCP, UDP)

Use of Port Numbers

- **Client applications chose a free port number (which is not already used by another connection) as the source port**
- **The destination port is the well-known port of the server application**
- **Some services like FTP or Remote Procedure Call use dynamically assigned port numbers:**
  – Sun RPC (Remote Procedure Call) uses a portmapper located at port 111...
  – FTP uses the PORT and PASV commands...
- **...to switch to a non-standard port**

Internet Transport Layer, v4.5 13

Well Known Ports

Some Well Known Ports

7	Echo		
20	FTP (Data), File Transfer Protocol	**Some Registered Ports**	
21	FTP (Control)	1416	Novell LU6.2
23	TELNET, Terminal Emulation	1433	Microsoft-SQL-Server
25	SMTP, Simple Mail Transfer Protocol	1439	Eicon X25/SNA Gateway
53	DNS, Domain Name Server	1527	oracle
69	TFTP, Trivial File Transfer Protocol	1986	cisco license managmt
80	HTTP Hypertext Transfer Protocol	1998	cisco X.25 service (XOT)
111	Sun Remote Procedure Call (RPC)	5060	SIP (VoIP Signaling)
137	NetBIOS Name Service	6000	\
138	NetBIOS Datagram Service	.....	> X Window System
139	NetBIOS Session Service	6063	/
161	SNMP, Simple Network Management Protocol		
162	SNMPTRAP	... etc.	
322	RTSP (Real Time Streaming Protocol) Server	(see RFC1700)	

Internet Transport Layer, v4.5 14

TCP Header

Internet Transport Layer, v4.5 15

TCP Header Entries

- **Source and Destination Port**
  – Port number for source and destination process
- **Header Length**
  – Indicates the length of the header given as a multiple of 32 bit words (4 octets)
  – necessary, because of the variable header length
- **Sequence Number (32 Bit)**
  – Position of the first octet of this segment within the data stream ("wraps around" to 0 after reaching $2^{32} - 1$)
- **Acknowledge Number (32 Bit)**
  – Acknowledges the correct reception of all octets up to ack-number minus 1 and indicates the number of the next octet expected by the receiver

Internet Transport Layer, v4.5 16

TCP Header Entries

- **Flags: SYN, ACK**

  – SYN: If set, the Sequence Number holds the initial value for a new session
    - SYN is used only during the connect phase (can be used to recognize who is the caller during a connection setup e.g. for firewall filtering)
  – Used for call setup (connect request)

  – ACK: If set, the Acknowledge Number is valid and indicates the sequence number of the next octet expected by the receiver

TCP Header Entries

- **Flags: FIN, RST**

  – FIN: If set, the Sequence Number holds the number of the last transmitted octet of this session
    - using this number a receiver can tell that all data have been received; FIN is used only during the disconnect phase

  – Used for call release (disconnect)

  – RST: If set, the session has to be cleared immediately
    - Can be used to refuse a connection-attempt or to "kill" a current connection.

TCP Header Entries

- **Window (16 Bit)**
  – Set by the source with every transmitted segment to signal the current window size; this "dynamic windowing" enables receiver-based flow control
  – The value defines how many additional octets will be accepted, starting from the current acknowledgment number
    - SeqNr of last octet allowed to sent: AckNr plus window value
  – Remarks:
    - Once a given range for sending data was given by a received window value, it is not possible to shrink the window size to such a value which gets in conflict with the already granted range
    - so the window field must be adapted accordingly in order to achieve the flow control mechanism STOP

TCP Header Entries

- **Checksum**
  – The checksum includes the TCP header and data area plus a 12 byte pseudo IP header
    - (one´s complement of the sum of all one´s complements of all 16 bit words)
  – The pseudo IP header contains the source and destination IP address, the IP protocol type and IP segment length (total length). This guarantees, that not only the port but the complete socket is included in the checksum
  – Including the pseudo IP header in the checksum allows the TCP layer to detect errors, which can't be recognized by IP (e.g. IP transmits an error-free TCP segment to the wrong IP end system)

TCP Header Entries

● **Flags: URG**

  – Is used to indicate important - "urgent" - data

  – If set, the 16-bit "Urgent Pointer" field is valid and points to the <u>last</u> octet of urgent data

    • sequence number of last urgent octet = actual segment sequence number + urgent pointer

    • RFC 793 and several implementations assume the urgent pointer to point to the first octet *after* the urgent data; However, the "Host Requirements" RFC 1122 states this as a mistake!

    • Note: There is no way to indicate the beginning of urgent data (!)

  – When a TCP receives a segment with the URG flag set, it notifies the application which switch into the "urgent mode" until the last octet of urgent data is reached

  – Examples for use: Interrupt key in Telnet, Rlogin, or FTP

TCP Header Entries

● **Urgent Pointer**

  – points to the last octet of urgent data

● **Options**

  – Only MSS (Maximum Segment Size) is used frequently Other options are defined in RFC1146, RFC1323 and RFC1693

● **Pad**

  – Is used to make the header length an integral number of 32 bits (4 octets) because of the variable length options

TCP Header Entries

● **Flags: PSH**

  – "PUSH": If set, the segment should be forwarded to the next layer immediately without buffering

    A TCP instance can decide on its own, when to send data to the next instance. One strategy could be, to collect data in a buffer and forward the data when the buffer exceeds a certain size. An application which needs a low latency or a constant data stream would like to bypass this buffer with the PSH flag. Also the last segment of a request could use the PSH flag.
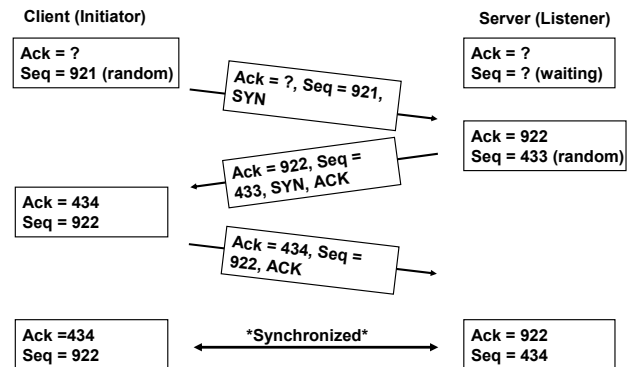
  – Today often ignored

TCP Connect Phase

TCP Connect Phase

- **TCP uses the unreliable service of IP, hence TCP segments of old sessions (e.g. retransmitted or delayed segments) could disturb a TCP connect**
- **Random starting sequence numbers and an explicit negotiation of starting sequence numbers makes a TCP connect immune against spurious packets**
- **Disturbing Segments (e.g. delayed TCP segments from old sessions etc.) and old "half-open" connections are deleted with the RST flag**
- **--> "Three Way Handshake"**

TCP Duplicates after Connect

- **Duplicates of old TCP sessions (same source/destination IP address and TCP socket) can disturb a new session**
- **Thus sequence numbers must be unique for different sessions of the same socket.**
- **Initial sequence number (ISN) must be chosen with a good algorithm**
- **RFC793 suggests to pick a random number at boot time (e.g. derived from system start up time) and increment every 4 μs. Every new connection will increment additionally by 1**
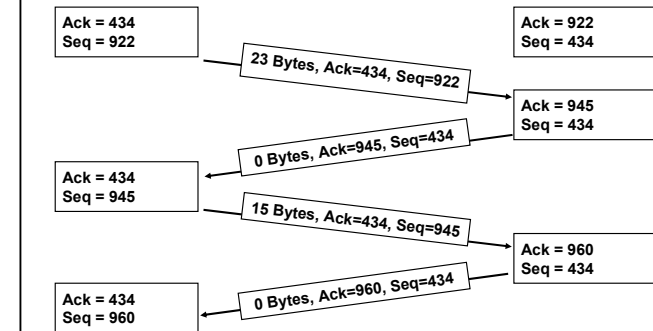
TCP Data Transfer Phase

TCP Data Transfer Phase

- **Acknowledgements are generated for all octets which arrived in sequence without errors (positive acknowledgement)**
  - Note: duplicates are also acknowledged
  - If a segment arrives out of sequence, no acknowledges are sent until this "gap" is closed
- **The acknowledge number is equal to the sequence number of the next octet to be received**
  - Acknowledges are "cumulative": Ack(N) confirms all octets with sequence numbers up to N-1
  - Thus, lost acknowledgements are not critical since the following ack confirms all previous segments
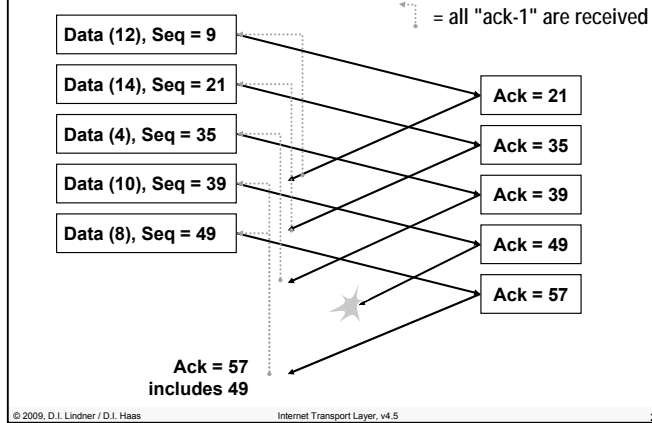
L11 - Internet Transport Layer (TCP, UDP)

TCP "Cumulative" Acknowledgement

= all "ack-1" are received

Data (12), Seq = 9
Data (14), Seq = 21
Data (4), Seq = 35
Data (10), Seq = 39
Data (8), Seq = 49

Ack = 21
Ack = 35
Ack = 39
Ack = 49
Ack = 57

Ack = 57 includes 49

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   29
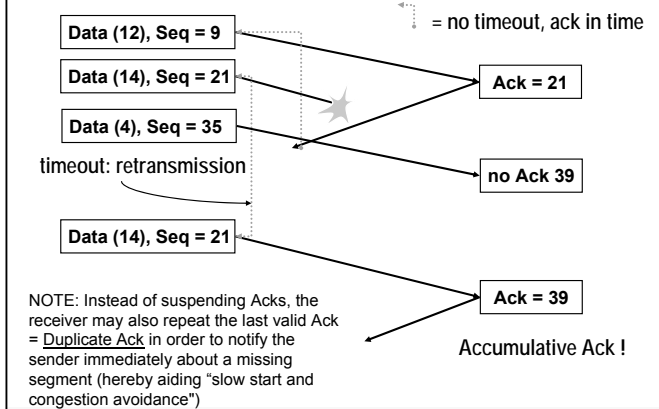

TCP Timeout

- **Timeout will initiate a retransmission of unacknowledged data**

- **Value of retransmission timeout influences performance (timeout should be in relation to round trip delay)**
  – High timeout results in long idle times if an error occurs
  – Low timeout results in unnecessary retransmissions

- **Adaptive timeout**
  – KARN algorithm uses a backoff method to adapt to the actual round trip delay

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   30
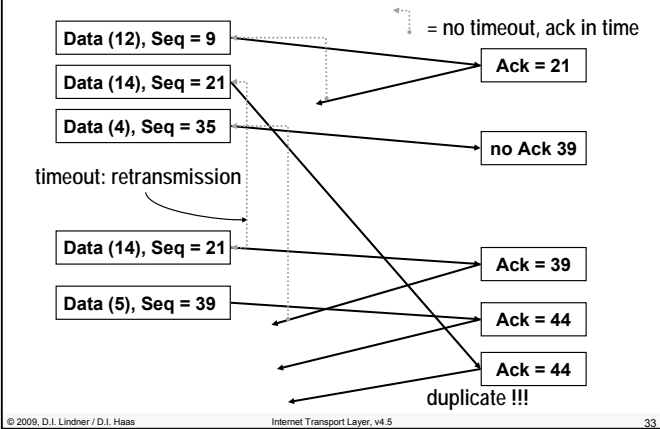

TCP Duplicates

- **Reasons for <u>retransmission</u>:**
  – Because original segment was lost:
    No problem, retransmitted segment fills gap, no duplicate
  – Because ACK was lost or retransmit timeout expired:
    No problem, segment is recognized as duplicate through the sequence number
  – Because original was delayed and timeout expired:
    No problem, segment is recognized as duplicate through the sequence number

- **32 bit sequence numbers provide enough "space" to differentiate duplicates from originals**
  – 2^{32} Octets with 2 Mbit/s means 9h for wrap around (compare to usual TTL = 64 seconds)

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   31

TCP Duplicates, Lost Original (old TCP)

= no timeout, ack in time

Data (12), Seq = 9
Data (14), Seq = 21
Data (4), Seq = 35

timeout: retransmission

Data (14), Seq = 21

Ack = 21
no Ack 39
Ack = 39

Accumulative Ack !

NOTE: Instead of suspending Acks, the receiver may also repeat the last valid Ack = <u>Duplicate Ack</u> in order to notify the sender immediately about a missing segment (hereby aiding "slow start and congestion avoidance")

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   32

© 2009, D.I. Lindner / D.I. Haas

© 2009, D.I. Lindner / D.I. Haas

L11 - Internet Transport Layer (TCP, UDP)

TCP Duplicates, Delayed Original

- Data (12), Seq = 9
- Data (14), Seq = 21
- Data (4), Seq = 35
- timeout: retransmission
- Data (14), Seq = 21
- Data (5), Seq = 39
- = no timeout, ack in time
- Ack = 21
- no Ack 39
- Ack = 39
- Ack = 44
- Ack = 44
- duplicate !!!

TCP Duplicates, Lost Acknowledgement

- Data (12), Seq = 9
- Data (14), Seq = 21
- timeout: retransmission
- Data (14), Seq = 21
- Ack = 21
- Ack = 35
- Ack = 35
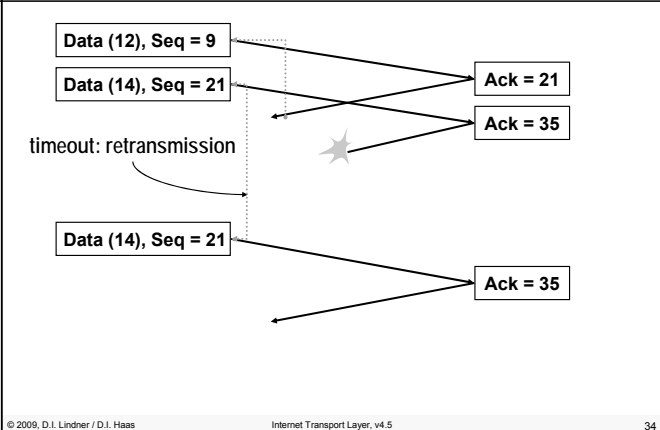
L11 - Internet Transport Layer (TCP, UDP)

TCP Duplicate-Ack (new TCP)

- Data (12), Seq = 9
- Data (14), Seq = 21
- Data (4), Seq = 35
- Data (10), Seq = 39
- timeout: retransmission
- Data (14), Seq = 21
- = no timeout, ack in time
- Ack = 21
- Duplicate Acks !
- Ack 21
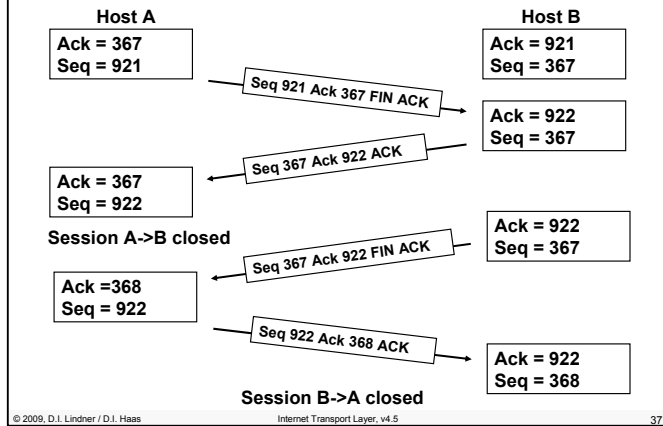- Ack 21
- Ack = 49
- Accumulative Ack !
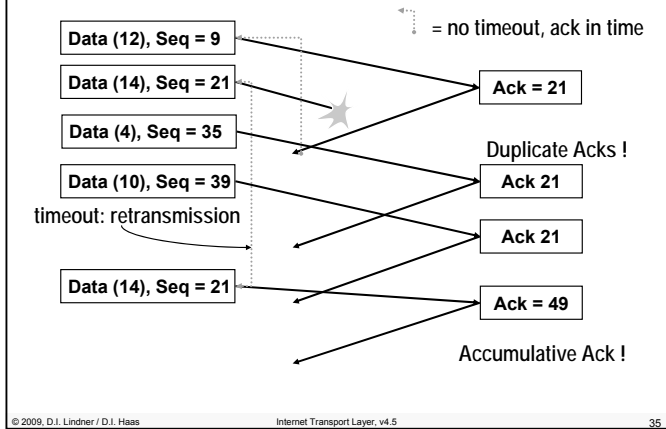
TCP Disconnect

- **A TCP session is disconnected similar to the three way handshake**
- **The FIN flag marks the sequence number to be the last one; the other station acknowledges and terminates the connection in this direction**
- **The exchange of FIN and ACK flags ensures, that both parties have received all octets**
- **The RST flag can be used if an error occurs during the disconnect phase**
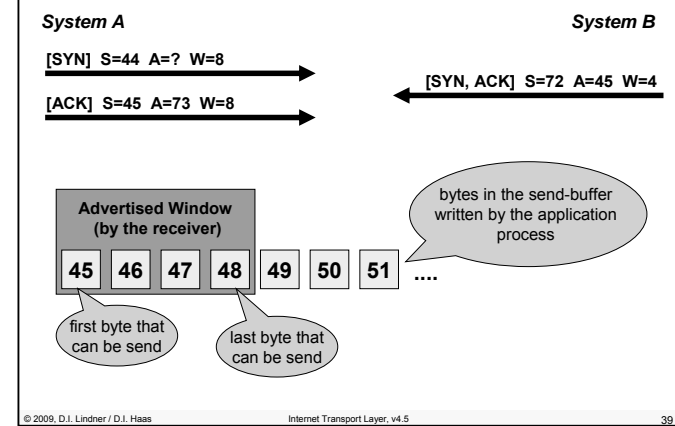
L11 - Internet Transport Layer (TCP, UDP)

TCP Disconnect

Host A

Ack = 367
Seq = 921

Seq 921 Ack 367 FIN ACK

Ack = 921
Seq = 367

Host B

Ack = 922
Seq = 367

Seq 367 Ack 922 ACK

Ack = 367
Seq = 922

Session A->B closed

Ack = 922
Seq = 367

Seq 367 Ack 922 FIN ACK

Ack =368
Seq = 922

Seq 922 Ack 368 ACK

Ack = 922
Seq = 368

Session B->A closed

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   37

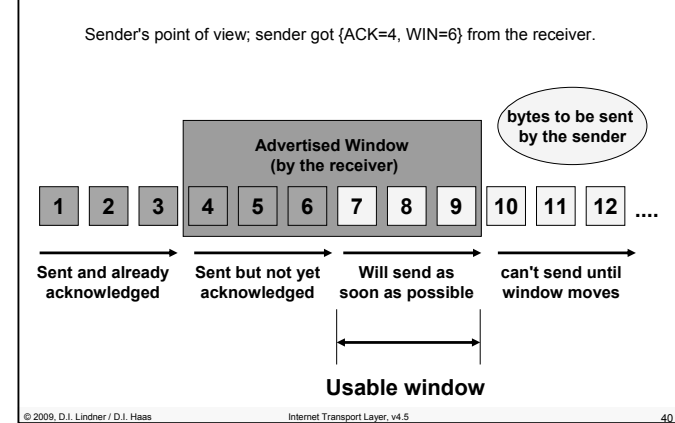## Flow control: "Sliding Window"

- **TCP flow control is done with dynamic windowing using the sliding window protocol**
- **The receiver advertises the current amount of octets it is able to receive**
  - using the window field of the TCP header
  - values 0 through 65535
- **Sequence number of the last octet a sender may send = received ack-number -1 + window size**
  - The starting size of the window is negotiated during the connect phase
  - The receiving process can influence the advertised window, hereby affecting the TCP performance

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   38

Sliding Window: Initialization

System A   *System B*

[SYN] S=44 A=? W=8

[SYN, ACK] S=72 A=45 W=4

[ACK] S=45 A=73 W=8

Advertised Window (by the receiver)

bytes in the send-buffer written by the application process

45 46 47 48 49 50 51 ....

first byte that can be send

last byte that can be send

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   39

Sliding Window: Principle

Sender's point of view; sender got {ACK=4, WIN=6} from the receiver.

bytes to be sent by the sender

Advertised Window (by the receiver)

1 2 3 4 5 6 7 8 9 10 11 12 ....

Sent and already acknowledged   **Sent but not yet acknowledged**   **Will send as soon as possible**   **can't send until window moves**

Usable window

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   40
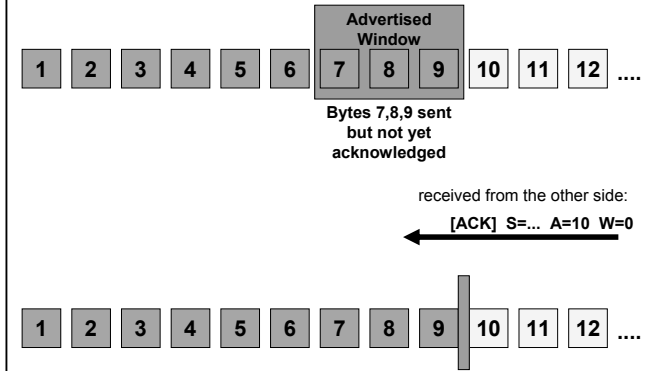
L11 - Internet Transport Layer (TCP, UDP)
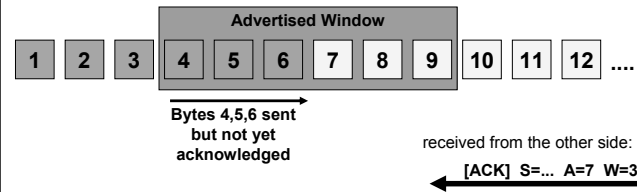
Sliding Window

- **During the transmission the sliding window moves from left to right, as the receiver acknowledges data**
- **The relative motion of the two ends of the window *open* or *closes* the window**
  - the window closes when data - already sent - is acknowledged (the left edge advances to the right)
  - the window opens when the receiving process on the other end reads data - and hence frees up TCP buffer space - and finally acknowledges data with a appropriate window value (the right edge moves to the right)
- **If the left edge reaches the right edge, the sender stops transmitting data - *zero usable window***

Internet Transport Layer, v4.5   41

L11 - Internet Transport Layer (TCP, UDP)

Flow Control -> STOP, Window Closed

Advertised Window

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

Bytes 7,8,9 sent but not yet acknowledged

received from the other side:

[ACK] S=... A=10 W=0

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

Internet Transport Layer, v4.5   43

Closing the Sliding Window

Advertised Window

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

Bytes 4,5,6 sent but not yet acknowledged

received from the other side:

[ACK] S=... A=7 W=3
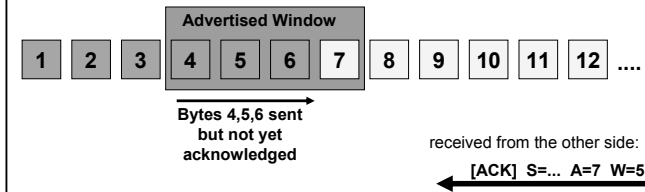
Advertised Window

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

Now the sender may send bytes 7, 8, 9. The receiver didn't open the window (W=3, right edge remains constant) because of congestion. However, the remaining three bytes inside the window are already granted, so the receiver cannot move the right edge leftwards.

Internet Transport Layer, v4.5   42

Opening the Sliding Window

Advertised Window

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

Bytes 4,5,6 sent but not yet acknowledged

received from the other side:

[ACK] S=... A=7 W=5

Advertised Window

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | .... |

The receiver's application read data from the receive-buffer and acknowledged bytes 4,5,6. Free space of the receiver's buffer is indicated by a window value that makes the right edge of the window move rightwards. Now the sender may send bytes 7, 8, 9,10,11.

Internet Transport Layer, v4.5   44

Sliding Window

- **The right edge of the window must not move leftward !**
  - however, TCP must be able to cope with a peer doing this
  - called *shrinking* window
- **The left edge cannot move leftward because it is determined by the acknowledgement number of the receiver**
  - only a duplicate Ack would imply to move the left edge leftwards, but duplicate Acks are discarded

TCP Enhancements

- **So far, only basic TCP procedures have been mentioned**
- **TCP's development still continues; it has been already enhanced with additional functions which are essential for operation of TCP sessions in today's IP networks:**
  - Slow Start and Congestion Avoidance Mechanism
  - Fast Retransmit and Fast Recovery Mechanism
  - Delayed Acknowledgements
  - The Nagle Algorithm
  - ....

TCP Enhancements Overview   1

- **Slow Start and Congestion Avoidance Mechanism:**
  - controls the rate of packets which are put into a network (sender-controlled flow control as add on to the receiver-controlled flow control based on the window field)

- **Fast Retransmit and Fast Recovery Mechanism:**
  - to avoid waiting for the timeout in case of retransmission and to avoid slow start after a fast retransmission

TCP Enhancements Overview   2

- **Immediate acknowledgements may cause an unnecessary amount of data transmissions**
  - normally, an acknowledgement would be send immediately after the receiving of data
  - but in interactive applications, the send-buffer at the receiver side gets filled by the application soon after an acknowledgement has been send (e.g. Telnet echoes)
- **In order to support piggy-backed acknowledgements (i.e. acks combined with user data), the TCP stack waits 200 ms before sending the <u>delayed acknowledgement</u>**
  - during this time, the application might also have data to send

TCP Enhancements Overview 3

- **Several applications send only very small segments - "tinygrams"**
  - e.g. Telnet or Rlogin where each key-press generates 41 bytes to be transmitted: 20 bytes IP header, 20 bytes TCP header and only 1 byte of data (!)
- **Frequent tinygrams can lead to congestion at slow WAN connections**
- **Nagle Algorithm:**
  - When a TCP connection waits for an acknowledgement, small segments must not be sent until the acknowledgement arrives
  - In the meanwhile, TCP can collect small amounts of application data and send them in a single segment when the acknowledgement arrives

Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - Slow Start and Congestion Avoidance
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- **RFC Collection**

Transport Layer Protocols
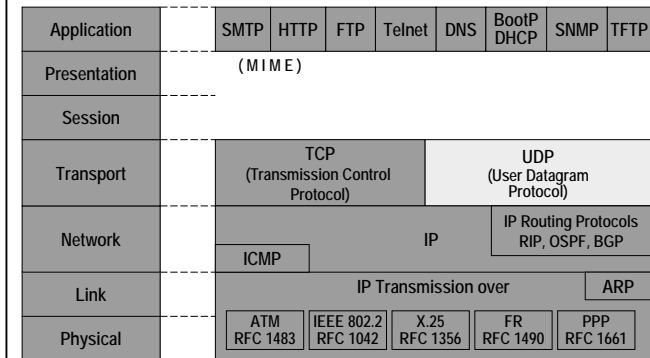
UDP (User Datagram Protocol, RFC 768)

- **UDP is a connectionless layer 4 service (datagram service)**
- **Layer 3 Functions are extended by port addressing and a checksum to ensure integrity**
- **UDP uses the same port numbers as TCP (if applicable)**
- **Less complex than TCP, easier to implement**

UDP and OSI Transport Layer 4

Layer 4 Protocol = UDP (Connectionless)

IP Host A   IP Host B

UDP Connection (Transport-Pipe)

↕ 4   4 ↕
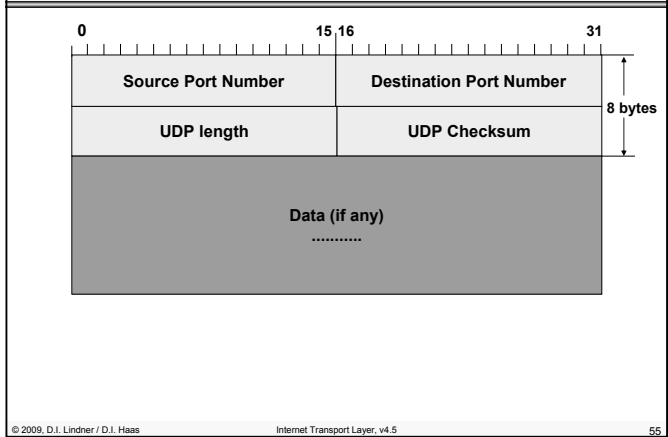
M   M

Router 1   Router 2
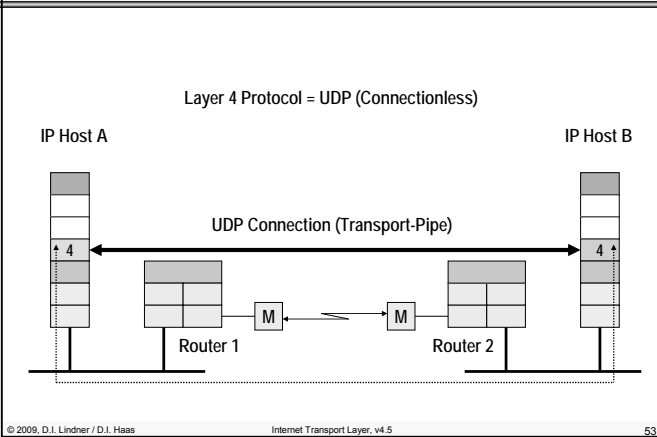
UDP Usage

- **UDP is used**
  - where the overhead of a connection oriented service is undesirable
    - e.g. for short DNS request/reply
  - where the implementation has to be small
    - e.g. BootP, TFTP, DHCP, SNMP
  - where retransmission of lost segments makes no sense
    - Voice over IP
    - note: digitized voice is critical concerning delay but not against loss
      - Voice is encapsulated in RTP (Real-time Transport Protocol)
      - RTP is encapsulated in UDP
      - RTCP (RTP Control Protocol) propagates control information in the opposite direction
      - RTCP is encapsulated in UDP

UDP Header

0	15	16	31
Source Port Number		Destination Port Number	
UDP length		UDP Checksum	

8 bytes

Data (if any)
....

UDP Header Entries

- **Source and Destination Port**
  - Port number for addressing the process (application)
  - Well known port numbers defined in RFC1700

- **UDP Length**
  - Length of the UDP datagram (Header plus Data)

- **UDP Checksum**
  - Checksum includes pseudo IP header (IP src/dst addr., protocol field), UDP header and user data. One´s complement of the sum of all one´s complements

L11 - Internet Transport Layer (TCP, UDP)

Important UDP Port Numbers

- – 7   Echo
- – 53   DOMAIN, Domain Name Server
- – 67   BOOTPS, Bootstrap Protocol Server
- – 68   BOOTPC, Bootstrap Protocol Client
- – 69   TFTP, Trivial File Transfer Protocol
- – 79   Finger
- – 111   SUN RPC, Sun Remote Procedure Call
- – 137   NetBIOS Name Service
- – 138   NetBIOS Datagram Service
- – 161   SNMP, Simple Network Management Protocol
- – 162   SNMP Trap
- – 322   RTSP (Real Time Streaming Protocol) Server
- – 520   RIP
- – 5060   SIP (VoIP Signaling)
- – xxxx   RTP (Real-time Transport Protocol)
- – xxxx+1 RTCP (RTP Control Protocol)

Elements and Scenarios of VoIP

Service Management Server

Gateway   Gateway

PSTN   Internet   PSTN
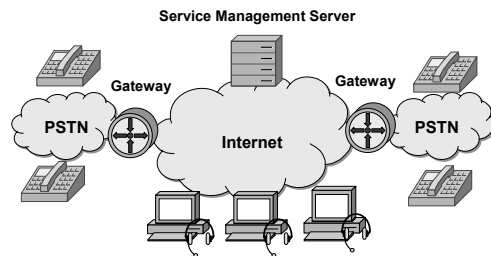
Elements: PCs, conventional telephones, IP-telephones
Scenarios: PC-to-PC, Phone-to-Phone, PC-to-Phone, Phone-to-PC
PSTN = public switched telephone network

L11 - Internet Transport Layer (TCP, UDP)

VoIP Clients, Gateways, Servers

- **Clients**
  - – User-interface
  - – To call or end a call
  - – Analogue – digital encoding of speech
  - – IP-data packetizing
  - – Decode from digital into analogue speech

  2 Types:
  - – Clients: software-clients or IP-telephones
  - – Virtual clients: provided by gateway, interface for conventional telecommunication equipment like telephones, fax etc.

- **VoIP-Gateway**
  - – Bridge between conventional and IP telephony
  - – Allows both users to communicate with their different equipment

- **Server**
  - – IP-telephony management and control
  - – Management of connection requirements of connection and exchange processes like
  - – Call forwarding, conference calls, user administration of their profiles and access rights, call tracking, billing, answering machine, voice-mail function

Tasks of a VoIP Gateway

- **Task 1: matching telephone number - IP-address**
  whenever you call from a conventional phone over the VoIP gateway, the server has to convert the wanted phone number into an IP-address of the remote gateway of the call receiver with a database lookup.
  - – peripheral database => directly implemented in gateway performance advantages in speed
  - – central database on a server where all gateways have access bigger latency but no database replication is needed

- **Task 2: Connection establishment**
  - – Gateway is contacting the remote-gateway and exchanging security, encoding, capacity and setup information until connection is established

Tasks of a VoIP Gateway

- **Task 3: digitalization and compression**
  - Analogue speech signals have to be digitized before compression.
  - Common technique: 64K Pulse Code Modulation (PCM)
  - ISDN channels can be easily connected because they are already 64K PCM encoded and can be bridged. Compression into one of several codec-formats is done by the Digital Signal Processor (DSP)

- **Task 4: Packetizing and packet delivery**
  - Wrap data into IP packets and dispatch them via UDP and TCP
  - Advantage of UDP: no error detection and recovery => faster, more efficient, retransmission of speech wouldn't make sense => delay

Current Problems of Internet Telephony   1

- **Standards**
  - interoperability between Internet telephony products and PSTN-based systems and services
  - Users have to have the same kind of software

- **Quality**
  - Voice performance is measured by delay
  - Calls on PSTN have about 50-70 msec delay
  - On internet there is an increased latency of ~ 500 msec
  - But human latency tolerance is only ~ 250 msec
  - Today's products exceed it so it sound like calls routed over a satellite circuit

Current Problems of Internet Telephony   2

- **Capacity**
  - Packet loss occur because of network congestion due to
    - bandwidth limitation
    - traffic overload → transmission delays and packet discards
    - Error performance → inadequate network access links cause bandwidth congestion (very bad on transcontinental links)
    - applications repair lost packets with silence → speech clipping effects → Even the loss of an individual packet has an impact on speech due to the large packet size.

- **Social issues**
  - Traditional telephone providers (often monopolies) are against Internet-based providers because they have an "unfair" advantage in offering cut-rate long distance phone service.
  - Conclusion: It is very political!

Solution to Current Problems   1

- **Standards**
  - ITU H.323 recommendation for VoN applications
  - Improved voice compression codecs
  - T.120 for data conferencing
  - RTP (Real Time Protocol)
  - RTCP (Real Time Control Protocol)
  - IP QoS (IntSrv with Resource Reservation Protocol or DiffSrv with DSCP)
  - SIP (Session Initiation Protocol)

- **Quality improvements**
  - protocol improvements (IP QoS) and codec improvements
  - bigger routers (gigabit routers)
  - new network architectures and better links

L11 - Internet Transport Layer (TCP, UDP)

Solution to Current Problems   2

- **Capacity**
  - average hop number of trans-Atlantic call is 20 to 30
  - delay increase with every router hop
    - increase routing speed
    - more routers
    - bigger routers (gigabit router)
      - handle at least 10 times more traffic than conventional router
      - per-packet cost of gigabit routing is 3-4 times less than traditional routing
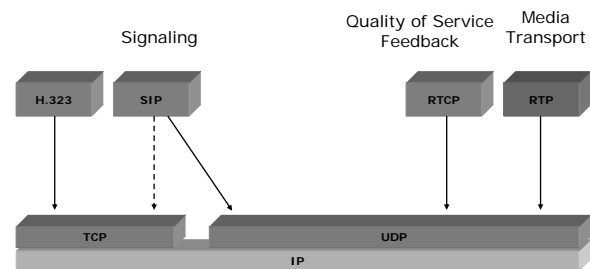
- **Social issues**
  - encourage new technology
  - trend of technology development
    - more than 100 well known companies are involved

VoIP Protocols - Overview

L11 - Internet Transport Layer (TCP, UDP)

Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - <u>Slow Start and Congestion Avoidance</u>
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- **RFC Collection**

How to Improve TCP Performance?

- **Problem:**
  - in case of packet loss sender can use the window given by the receiver but when window becomes closed the sender must wait until <u>retransmission timer</u> times out
  - that means during that time sender cannot use offered bandwidth of the network
  - -> TCP performance degradation
- **Assumption:**
  - packet loss in today´s networks are mainly caused by congestion but not by bit errors on physical lines
    - optical transmission
    - digital transmission

Congestion

- **Problem of bottle-neck inside of a network**
  - Some intermediate router must queue packets
  - Queue overflow -> retransmission -> more overflow !
  - Can't be solved by traditional receiver-imposed flow control (using the window field)
- **Ideal case: rate at which new segments are injected into the network = acknowledgment-rate of the other end**
  - Requires a sensitive algorithm to catch the equilibrium point between high data throughput and packet dropping due to queue overflow:
    Van Jacobson's Slow Start and Congestion Avoidance

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   69

Slow Start   1

- **Duplicate Ack can be used for catching this equilibrium**
  - That is the base for Slow Start and Congestion Avoidance

- **Slow start (and congestion avoidance) is mandatory for today's TCP implementations !**

- **Slow start requires TCP to maintain an additional window: the "congestion window" (cwnd)**
  - Rule: *The sender may transmit up to the minimum of the congestion window and the advertised window*

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   70
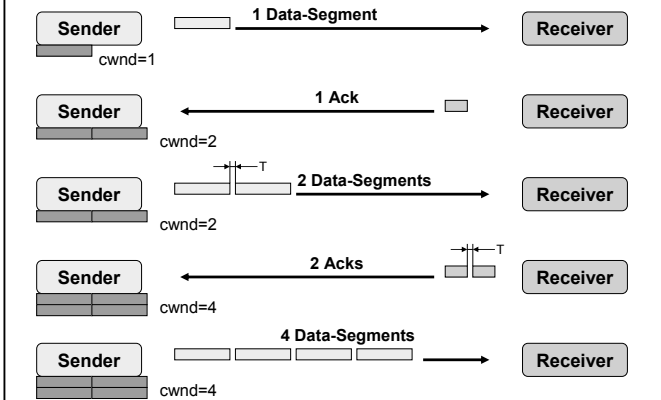
Slow Start   2

- **When a new TCP connection is established, the congestion window is initialized to <u>one segment</u>**
  - Using the maximum segment size (MSS) of the receiver (learned via a TCP option field)
    - Note: RFC 2414 suggests increasing the initial congestion window to 2-4 segments
- **Each time the sender receives an acknowledgment, the congestion window is increased by one segment size**
- **This way, the segment send rate <u>doubles</u> every round trip time until congestion occurs; then the sender has to slow down again**

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   71

Slow Start   3

© 2009, D.I. Lindner / D.I. Haas   Internet Transport Layer, v4.5   72

L11 - Internet Transport Layer (TCP, UDP)

Congestion

- **Slow start encounters congestion, when**
  - routers connect pipes with different bandwidth
  - routers combine several input pipes and the bandwidth of the output pipe is less than the sum of the input bandwidths
  - and hence TCP segment(s) is (are) dropped by a router
- **Congestion can be detected by the sender through <u>timeouts</u> or <u>duplicate acknowledgements</u>**
- **Slow start reduces its sending rate with the help of another algorithm, called "<u>Congestion Avoidance</u>"**

Congestion Avoidance   1

- **Slow start with congestion avoidance is a <u>sender-imposed flow control</u>**
  - Congestion Avoidance requires TCP to maintain a variable called "<u>slow start threshold</u>" (ssthresh)
  - Initially, ssthresh is set to TCPs maximum possible MSS (i.e. 65,535 octets)
- **On detection of congestion, ssthresh is set to <u>half</u> <u>the current window size</u>**
  - here, window size means: minimum of advertised window and congestion window (but at least 2 segments)
  - Note: ssthresh marks a safe window size because congestion occurred at a window size of 2 x ssthresh

L11 - Internet Transport Layer (TCP, UDP)

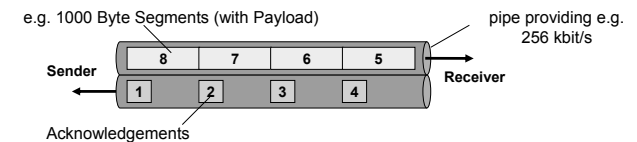## Congestion Avoidance   2

- **If the congestion is indicated by**
  - a timeout:
    - cwnd is set to 1 -> forcing slow start again
  - a duplicate ack:
    - cwnd is set to ssthresh (= 1/2 current window size)
- **cwnd ≤ ssthresh:**
  - slow start, doubling cwnd every round-trip time
  - <u>exponential</u> growth of cwnd
- **cwnd > ssthresh:**
  - congestion avoidance, cwnd is incremented by <u>MSS × MSS / cwnd</u> every time an ack is received
  - <u>linear</u> growth of cwnd

Slow Start and Congestion Avoidance
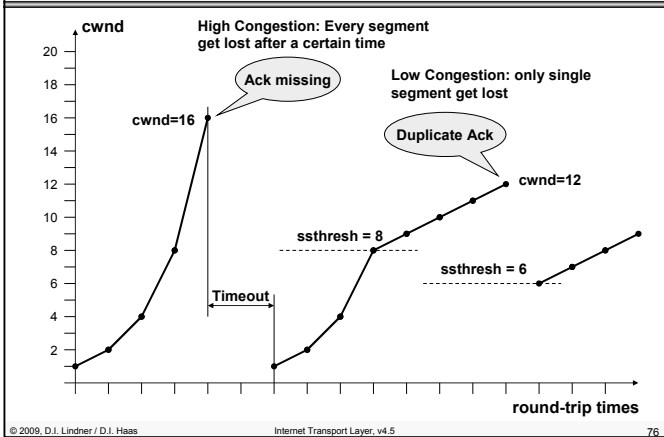
Long Term View of TCP Throughput
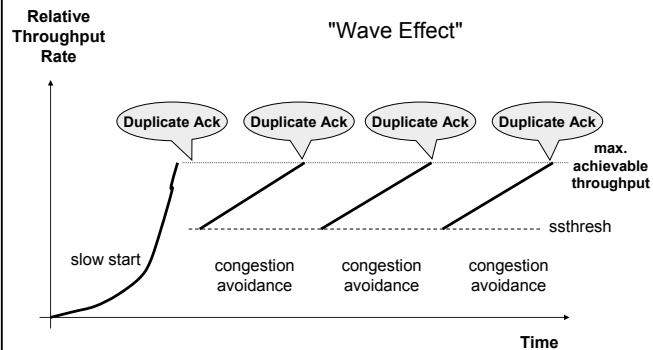
Limitation of ARQ Protocols

- **The performance of any connection-oriented protocol with error-recovery (ARQ) is limited by bandwidth and delay by nature!**
  - Optimum can be achieved by using Continous RQ with sliding window technique where the window is large enough to avoid stopping of sending
  - Large enough means to cover the time of the serialization and propagation delays
  - Note: senders and receivers window size maybe also be limited because of memory constraints

Limitation of ARQ Protocols

- **The sender's window must be big enough so that the sender can fully utilize the channel volume**

- **Channel volume is increased**
  - by delays caused by buffers
  - limited signal speed
  - Bandwidth

- **The channel volume can be expressed by the**
  Delay-Bandwidth Product

The Delay-Bandwidth Product

- **In order to fill the "pipe" between sender and receiver with data packets:**
  - the window-size advertised by the receiver should be not smaller than the Delay-Bandwidth Product of the pipe
  - window size $\geq$ capacity of pipe (bits)
    = bandwidth (bits/sec) x round-trip time (sec)

e.g. 1000 Byte Segments (with Payload)

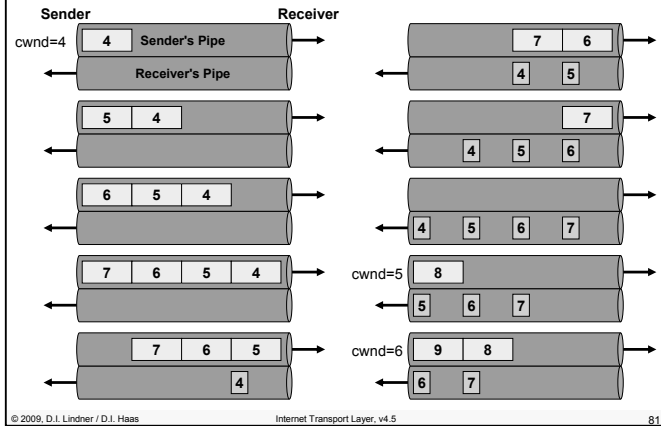pipe providing e.g. 256 kbit/s

Acknowledgements

Example: For a given RTT = 0.25 s (Round-trip time = elapsed time between the sending of segment n and the receiving of the corresponding ack n) the minimum window size is 256 kbit/s x 0.25 sec = 64 kbit. Using a segment size of 1 kB, the sender can transmit 8 segments before waiting for any acknowledgement.

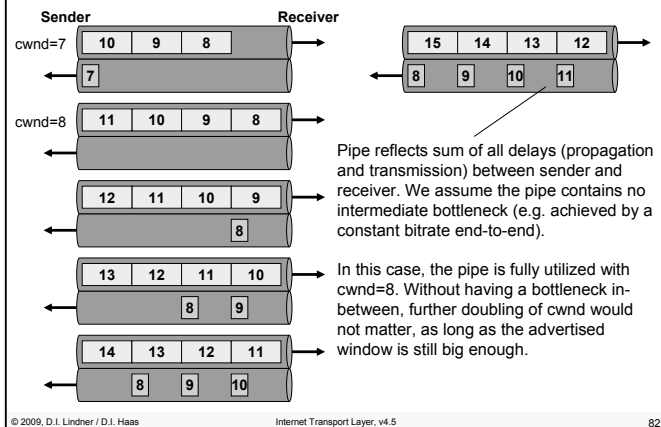**L11 - Internet Transport Layer (TCP, UDP)**

Filling the Pipe (1/2)

Filling the Pipe (2/2)

Pipe reflects sum of all delays (propagation and transmission) between sender and receiver. We assume the pipe contains no intermediate bottleneck (e.g. achieved by a constant bitrate end-to-end).

In this case, the pipe is fully utilized with cwnd=8. Without having a bottleneck in-between, further doubling of cwnd would not matter, as long as the advertised window is still big enough.

L11 - Internet Transport Layer (TCP, UDP)

Delay-Bandwidth Relations

Given pipe with given RTT and bandwidth:
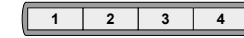
1) Doubled bandwidth:

Additional capacity
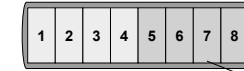
2) Doubled RTT:

Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - Slow Start and Congestion Avoidance
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- **RFC Collection**

Receiver Requirements

- **Initially, TCP could detect packet loss only by expiration of the retransmission timer**
  - receiver stops sending Acks until the sender retransmits all missing segments
  - causes long delays

- **Fast Retransmit requires a receiver to send an immediate <u>duplicate</u> acknowledgement in order to notify the sender which segments are (still) expected by the receiver**

Sender Aspects and Fast Retransmit

- **When should retransmission occur?**
  - Note: the receiver will also send duplicate acknowledgements when segments are arriving in the <u>wrong order</u>
  - Typically reordering problems cause one or two duplicate acks on the average

- **Therefore, TCP sender awaits two duplicate acknowledgements and starts retransmission <u>after the third</u> duplicate acknowledgement**
  - that mechanism is called "Fast Retransmit"

Fast Recovery   1

- **"Fast Retransmit" automatically commences "Fast Recovery" in order to rapidly repair single packet loss**

- **"Fast Recovery" mechanism:**
  - ssthresh is set to half the current window size
  - cwnd is set to ssthresh <u>plus 3 times the segment size</u>
    - Remember:
      Fast Retransmit waits for 3 duplicate acks; from this can be concluded that the receiver must have received 3 segments already

Fast Recovery   2

- **"Fast Recovery" mechanism cont.**
  - congestion avoidance, but not slow start is performed
    - Remember:
      The receiver can only generate a duplicate ack when another segment is received. That is: there are <u>still segments flowing</u> through the network! Slow start would reduce this flow abruptly!

  - For each additional duplicate ack, the sender increases cwnd by 1 segment size

  - Upon receiving an ack that acknowledges new data
    - cwnd is set to ssthresh
    - sender resumes normal congestion avoidance mode

L11 - Internet Transport Layer (TCP, UDP)

Fast Recovery   3

- **Fast Recovery**
  - allows the sender to continue to maintain the ack-clocked data rate for new data while the packet loss repair is being undertaken
   - note: if send window would be closed abruptly the synchronization via duplicate acks would be lost

  - still the single packet loss indicates congestion and back off to normal congestion avoidance mode must be done

Fast Retransmit and Fast Recovery

L11 - Internet Transport Layer (TCP, UDP)
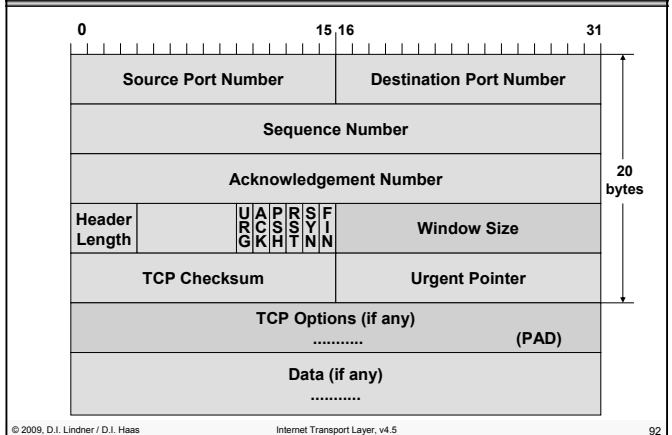
Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - Slow Start and Congestion Avoidance
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- **RFC Collection**
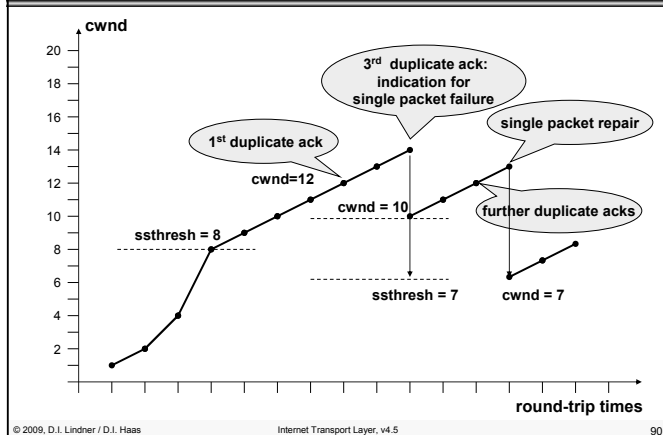
TCP Header Window Field

L11 - Internet Transport Layer (TCP, UDP)

TCP Options

- **Window-scale option**
  - a maximum segment size of 65,535 octets is inefficient for high delay-bandwidth paths
  - the window-scale option allows the advertised window size to be left-shifted (i.e. multiplication by 2)
  - <u>enables a maximum window size of 2^30 octets</u> !
  - negotiated during connection establishment

- **SACK (Selective Acknowledgement)**
  - if the SACK-permitted option is set during connection establishment, the receiver may selectively acknowledge already received data even if there is a gap in the TCP stream (Ack-based synchronization maintained)

Agenda

- **TCP Fundamentals**
- **UDP**
- **TCP Performance**
  - Slow Start and Congestion Avoidance
  - Fast Retransmit and Fast Recovery
  - TCP Window Scale Option and SACK Options
- <u>**RFC Collection**</u>

L11 - Internet Transport Layer (TCP, UDP)

RFCs

- **0761 - TCP**
- **0813 - Window and Acknowledgement Strategy in TCP**
- **0879 - The TCP Maximum Segment Size**
- **0896 - Congestion Control in TCP/IP Internetworks**
- **1072 - TCP Extension for Long-Delay Paths**
- **1106 - TCP Big Window and Nak Options**
- **1110 - Problems with Big Window**
- **1122 - Requirements for Internet Hosts -- Com. Layer**
- **1185 - TCP Extension for High-Speed Paths**
- **1323 - High Performance Extensions (Window Scale)**

RFCs

- **2001 - Slow Start and Congestion Avoidance (Obsolete)**
- **2018 - TCP Selective Acknowledgement (SACK)**
- **2147 - TCP and UDP over IPv6 Jumbograms**
- **2414 - Increasing TCP's Initial Window**
- **2581 - TCP Slow Start and Congestion Avoidance (Current)**
- **2873 - TCP Processing of the IPv4 Precedence Field**
- **3168 - TCP Explicit Congestion Notification (ECN)**