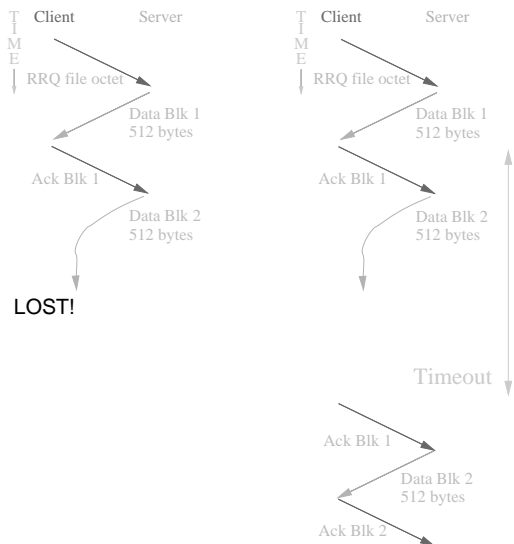


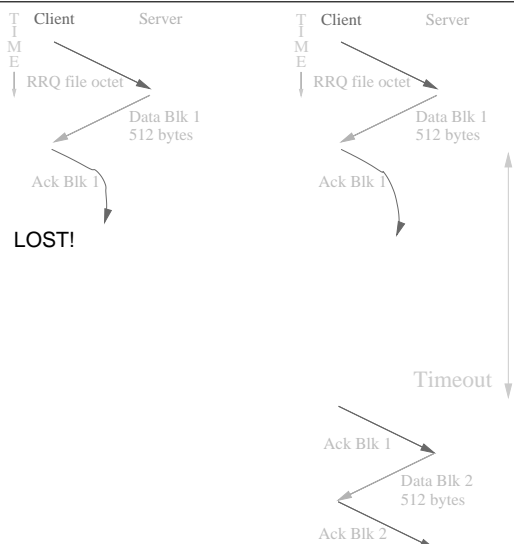
TFTP - lost packets?

- ◊ Packet recovery
- ◊ Timeout and retransmit
- ◊ Block number in data packets
 - and ACK
- ◊ Good enough.

Lost packet retransmit



Lost ACK?



TFTP - corrupted packets


- ◊ Data corruption
 - ◊ No TFTP checksum
 - ◊ Relies on UDP
 - ◊ and below
 - ◊ UDP checksum optional
 - ◊ Can have no checks
 - ◊ Not all that good
 - ◊ But just a bootstrap protocol

TFTP - duplications

- ◊ Packets delivered twice
- ◊ Several cases
 - ◊ RRQ / WRQ
 - ◊ Data
 - ◊ ACK
 - ◊ Error

TFTP Duplicate DATA

- ◊ Data packets carry block number
 - ◊ duplicate detection easy
- ◊ duplicates expected (loss recovery)



TFTP Duplicate ACK

- ◊ ACK is never duplicate
 - Always treated as ACK
 - Always transmit next block
 - (if any)
- ◊ Required for lost data packet recovery
 - May cause duplicate data packet
 - That is handled

TFTP Duplicate requests

- ◊ RRQ (or WRQ) is duplicated
- ◊ Each RRQ (WRQ) initiates a new connection
 - Server replies to each
 - from a different port number
- ◊ Client sees 2 replies
 - Knows it sent only one request
 - picks one reply - ignores the other

TFTP - out of order

- ◊ Packets delivered in wrong order
 - rare - only one outstanding packet
 - block number in data and ACK
- ◊ good enough

TFTP simultaneous transfers

- ◊ Multiple requests from one client
- ◊ Multiple requests to one server

- ◊ Port numbers differ
 - client picks a port number unique in its system
 - server sends from a port number unique in its system

- ◊ Together with IP address
 - those port numbers allow many requests

TFTP Large file transfer

- ◊ Block number in each packet
 - Block number is 16 bits

- ◊ Max block number is 65535
 - Each block 512 bytes

- ◊ Biggest file $65535 * 0.5\text{KB}$
 - just under 32 MB.

TFTP Throughput

- ◊ Work out maximum possible throughput of TFTP
 - Assume 100 Mbps ethernet & routers
 - Assume 2 ms RTT through network
 - Assume no delays in hosts

- ◊ 2ms RTT
- ◊ Lock Step protocol
 - 500 round trips per second

- ◊ 500 round trips, 512 bytes each time
 - 256 K Bytes/second

Increasing throughput

- ◊ Lock step protocol cannot work
 - we often cannot reduce the RTT
- ◊ So, need to allow multiple packets per RTT
 - send many packets before ACK of first
- ◊ Can send so much
 - that network cannot take more
 - (has its own problems)
 - need flow control
- ◊ TCP works this way

UDP

◊ User Datagram Protocol


- RFC 768
 - August 1980

This User Datagram Protocol (UDP) is defined to make available a datagram mode of packet-switched computer communication in the environment of an interconnected set of computer networks.

This protocol provides a procedure for application programs to send messages to other programs with a minimum of protocol mechanism.


The protocol is transaction oriented, and delivery and duplicate protection are not guaranteed.

User Datagram Protocol (UDP)




Source Port is an optional field, when meaningful, it indicates the port of the sending process, and may be assumed to be the port to which a reply should be addressed in the absence of any other information. If not used, a value of zero is inserted.

User Datagram Protocol (UDP)




Destination Port has a meaning within the context of a particular internet destination address.

User Datagram Protocol (UDP)



Length is the length in octets of this user datagram including this header and the data. (This means the minimum value of the length is eight.)

User Datagram Protocol (UDP)



Checksum is the 16-bit one's complement of the one's complement sum of a pseudo header of information from the IP header, the UDP header, and the data, padded with zero octets at the end (if necessary) to make a multiple of two octets.

Pseudo Header

- ◊ Pseudo == False (Fake, Pretend)

Source IP Address		
Destination IP Address		
0	Protocol	Payload Length

- ◊ Values taken from IP header
- ◊ Similar version defined for IPv6 (128 bit addresses)

Checksum Algorithm

- ◊ Inverse of 1's complement sum
 - of 16 bit words in data checksummed
- ◊ If result is 0, use -0 instead (0xFFFF)
 - $1AE3 + 34C2 ==> 4FA5$
 - $4FA5 + C207 ==> 11AD$
(2's complement, unlimited size: $4FA5+C207==111AC$)
 - Overflow is ignored,
 - but every time overflow occurs,
 - we have been past both -0 and 0
ie: have added 2
result is add of 1
so need to add an extra 1
 - Easy way is to
 - just add in the "overflow" value

UDP Checksums

- ◊ UDP Checksum is optional
 - Value of 0 indicates no checksum present
 - if checksum value had been 0
 - it would be represented as -0 (or FFFF)
 - Intended for applications where
 - data integrity is not important
 - But most link layers checksum packets,
 - and drop those with errors
 - Omitting checksum doesn't help much
 - And checksum also includes port numbers
 - packet could be delivered to wrong application,
 - or reply sent to wrong place
- ◊ Hence UDP checksum
 - now recommended to always be used.
 - mandatory in IPv6

UDP Evaluation

- ◊ Does the protocol work?
 - Lost packets?
 - Corrupted packets?
 - Out of order packets?
 - Duplicated packets?

- Multiple simultaneous transfers?

- ◊ It promises very little.

Introduction to TCP

- ◊ RFC 793 (+ RFC1122) (+ others)

- ◊ Transmission Control Protocol

The Transmission Control Protocol (TCP) is intended for use as a highly reliable host-to-host protocol between hosts in packet-switched computer communication networks, and in interconnected systems of such networks.

TCP is a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications.

TCP Header - Sequence Number

Source Port		Destination Port							
Sequence Number									
Acknowledgment Number									
HdrLen	Reserved	U	A	P	R	S	F	Window Size	
Checksum				Urgent Pointer					

- ◊ 32 bit number
- ◊ Every byte transmitted is allocated a number
- ◊ Sequence numbers cycle around

TCP Specification

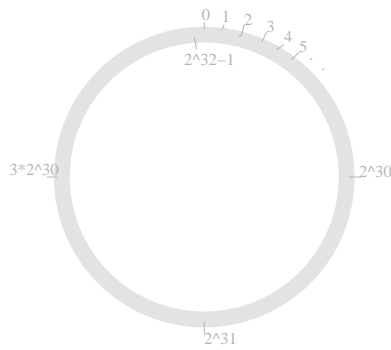
The TCP must recover from data that is damaged, lost, duplicated, or delivered out of order by the internet communication system.

This is achieved by assigning a sequence number to each octet transmitted, and requiring a positive acknowledgment (ACK) from the receiving TCP.

If the ACK is not received within a timeout interval, the data is retransmitted.

At the receiver, the sequence numbers are used to correctly order segments that may be received out of order and to eliminate duplicates.

TCP Sequence Number Space



◇ Sequence numbers run from 0 to $2^{32} - 1$

◇ 0, 1, 2, 3, ..., $2^{32}-2$, $2^{32}-1$, 0, 1, 2, ...

TCP Sequenced Data

Source Port				Destination Port															
Sequence Number																			
Acknowledgment Number																			
HdrLen		Reserved		U		A		P		R		S		F		Window Size			
Checksum								Urgent Pointer											
Data				Data															

◇ Sequence number in header identifies the first data byte in the packet

◇ Sequence numbers of later data bytes obtained by arithmetic