

CSC358 Intro. to Computer Networks

Lecture 5: Review, Transport Layer, (de)multiplexing, UDP, reliable data transfer

Amir H. Chinaei, Winter 2016

ahchinaei@cs.toronto.edu
http://www.cs.toronto.edu/~ahchinaei/

Many slides are (inspired/adapted) from the above source
© all material copyright; all rights reserved for the authors



Office Hours: T 17:00–18:00 R 9:00–10:00 BA4222

TA Office Hours: W 16:00–17:00 BA3201 R 10:00–11:00 BA7172
csc358ta@cdf.toronto.edu
http://www.cs.toronto.edu/~ahchinaei/teaching/2016jan/csc358/

Review


- ❖ Examples of Application Layer Protocols:
 - TLS, DNS, µTP
 - Proprietary protocols: e.g. Skype or myGame
- ❖ Examples of Transport Layer Protocols:
 - TCP, UDP, DCCP
- ❖ Examples of Network Layer Protocols:
 - IP, ICMP

- ❖ DNS (TCP/UDP)
- ❖ P2P

Application Layer 1-2

Client-server vs. P2P: example


client upload rate = u , $F/u = 1$ hour, $u_s = 10u$, $d_{min} \geq u_s$



Application Layer 2-3

P2P file distribution: BitTorrent


- ❖ file divided into 256Kb chunks
 - ❖ peers in torrent send/receive file chunks
- tracker*: tracks peers participating in torrent
torrent: group of peers exchanging chunks of a file



Application Layer 2-4

P2P file distribution: BitTorrent


- ❖ peer joining torrent:
 - has no chunks, but will accumulate them over time from other peers
 - registers with tracker to get list of peers, connects to subset of peers ("neighbors")
- ❖ while downloading, peer uploads chunks to other peers
- ❖ peer may change peers with whom it exchanges chunks
- ❖ *churn*: peers may come and go
- ❖ once peer has entire file, it may (selfishly) leave or (altruistically) remain in torrent



Application Layer 2-5

BitTorrent: requesting, sending file chunks


- requesting chunks*:
- ❖ at any given time, different peers have different subsets of file chunks
 - ❖ periodically, Alice asks each peer for list of chunks that they have
 - ❖ Alice requests missing chunks from peers, rarest first
- sending chunks: tit-for-tat*
- ❖ Alice sends chunks to those four peers currently sending her chunks at highest rate
 - other peers are choked by Alice (do not receive chunks from her)
 - re-evaluate top 4 every 10 secs
 - ❖ every 30 secs: randomly select another peer, starts sending chunks
 - "optimistically unchoke" this peer
 - newly chosen peer may join top 4



Application Layer 2-6

BitTorrent: tit-for-tat

- (1) Alice "optimistically unchokes" Bob
- (2) Alice becomes one of Bob's top-four providers; Bob reciprocates
- (3) Bob becomes one of Alice's top-four providers



Application Layer 2-7

Distributed Hash Table (DHT)

- ❖ Hash table
- ❖ DHT paradigm
- ❖ Circular DHT and overlay networks
- ❖ Peer churn

Application Layer 2-8

Simple Database

Simple database with (key, value) pairs:

- key: human name; value: social security #

Key	Value
John Washington	132-54-3570
Diana Louise Jones	761-55-3791
Xiaoming Liu	385-41-0902
Rakesh Gopal	441-89-1956
Linda Cohen	217-66-5609
.....
Lisa Kobayashi	177-23-0199

- key: movie title; value: IP address

Application Layer 2-9

Hash Table

- More convenient to store and search on numerical representation of key
- key = hash(original key)

Original Key	Key	Value
John Washington	8962458	132-54-3570
Diana Louise Jones	7800356	761-55-3791
Xiaoming Liu	1567109	385-41-0902
Rakesh Gopal	2360012	441-89-1956
Linda Cohen	5430938	217-66-5609
.....
Lisa Kobayashi	9290124	177-23-0199

Application Layer 2-10

Distributed Hash Table (DHT)

- ❖ Distribute (key, value) pairs over millions of peers
 - pairs are evenly distributed over peers
- ❖ Any peer can query database with a key
 - database returns value for the key
 - To resolve query, small number of messages exchanged among peers
- ❖ Each peer only knows about a small number of other peers
- ❖ Robust to peers coming and going (churn)

Application Layer 2-11



Assign key-value pairs to peers

- ❖ rule: assign key-value pair to the peer that has the *closest* ID.
- ❖ convention: closest is the *immediate successor* of the key.
- ❖ e.g., ID space $\{0, 1, 2, 3, \dots, 63\}$
- ❖ suppose 8 peers: 1, 12, 13, 25, 32, 40, 48, 60
 - If key = 51, then assigned to peer 60
 - If key = 60, then assigned to peer 60
 - If key = 61, then assigned to peer 1


Application Layer 2-12

Circular DHT

- each peer *only* aware of immediate successor and predecessor.




Resolving a query



$O(N)$ messages on average to resolve query, when there are N peers

Application Layer 2-14


Circular DHT with shortcuts



- each peer keeps track of IP addresses of predecessor, successor, short cuts.
- reduced from 6 to 3 messages.
- possible to design shortcuts with $O(\log N)$ neighbors, $O(\log N)$ messages in query

Application Layer 2-15

Peer churn




example: peer 5 abruptly leaves

handling peer churn:

- ❖ peers may come and go (churn)
- ❖ each peer knows address of its two successors
- ❖ each peer periodically pings its two successors to check aliveness
- ❖ if immediate successor leaves, choose next successor as new immediate successor

Application Layer 2-16

Peer churn



example: peer 5 abruptly leaves

- ❖ peer 4 detects peer 5's departure; makes 8 its immediate successor
- ❖ 4 asks 8 who its immediate successor is; makes 8's immediate successor its second successor.

handling peer churn:

- ❖ peers may come and go (churn)
- ❖ each peer knows address of its two successors
- ❖ each peer periodically pings its two successors to check aliveness
- ❖ if immediate successor leaves, choose next successor as new immediate successor

Application Layer 2-17


Let's move on to Transport Layer

- TCP, UDP
- principles, services
- multiplexing, demultiplexing
- reliable data transfer
- flow control
- congestion control

Transport Layer 3-18

Transport services and protocols

- ❖ provide *logical communication* between app processes running on different hosts
- ❖ transport protocols run in hosts
 - send side: breaks app messages into *segments*, passes to network layer
 - rcv side: reassembles segments into messages, passes to app layer
- ❖ more than one transport protocol available to apps
 - E.g.: TCP, UDP



Transport Layer 3-19

Transport vs. network layer

- ❖ **transport layer:** logical communication between A. processes
 - relies on, enhances, network layer services
- ❖ **network layer:** logical communication between T. processes
 - .relies on lower layer

household analogy:


12 kids in Ann's house sending letters to 12 kids in Bill's house:

- ❖ hosts = houses
- ❖ processes = kids
- ❖ app messages = (continuous) letter
- ❖ segments = letter (piece) in envelopes
- ❖ transport protocol = Ann and Bill who demux to in-house siblings
- ❖ network-layer protocol = postal service

Transport Layer 3-20

Internet transport-layer protocols

- ❖ reliable, in-order delivery (TCP)
 - congestion control
 - flow control
 - connection setup
- ❖ unreliable, unordered delivery: UDP
 - no-frills extension of "best-effort" IP
- ❖ services not available:
 - delay guarantees
 - bandwidth guarantees




Transport Layer 3-21

Multiplexing/demultiplexing

multiplexing at sender:
handle data from multiple sockets, add transport header (later used for demultiplexing)


demultiplexing at receiver:
use header info to deliver received segments to correct socket



Transport Layer 3-22

How demultiplexing works

- ❖ host receives IP datagrams
 - each datagram has source IP address, destination IP address
 - each datagram carries one transport-layer segment
 - each segment has source, destination port number
- ❖ host uses *IP addresses & port numbers* to direct segment to appropriate socket



TCP/UDP segment format

Transport Layer 3-23

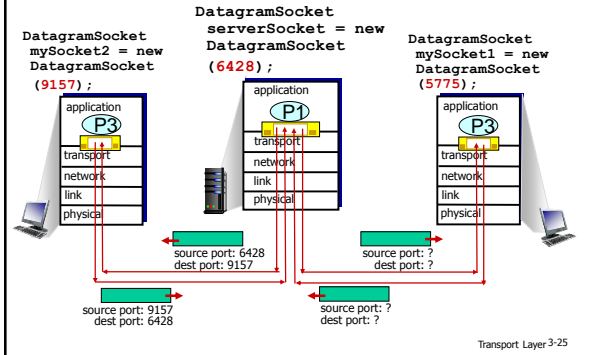
Connectionless demultiplexing

- ❖ **recall:** created socket has host-local port #:


```
DatagramSocket mySocket1 = new DatagramSocket(12534);
```
 - ❖ **recall:** when creating datagram to send into UDP socket, must specify
 - destination IP address
 - destination port #
 - ❖ when host receives UDP segment:
 - checks destination port # in segment
 - directs UDP segment to socket with that port #
- IP datagrams with *same dest. port #*, but different source IP addresses and/or source port numbers will be directed to *same socket* at dest

Transport Layer 3-24

Connectionless demux: example

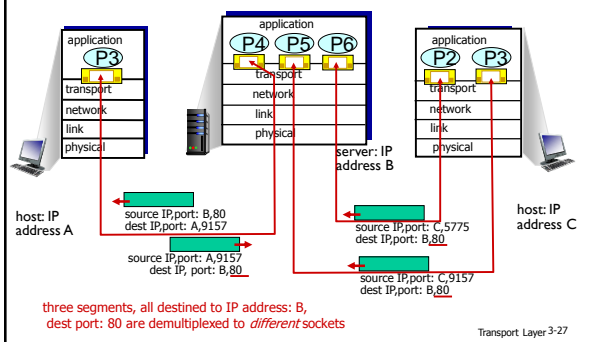


Connection-oriented demux

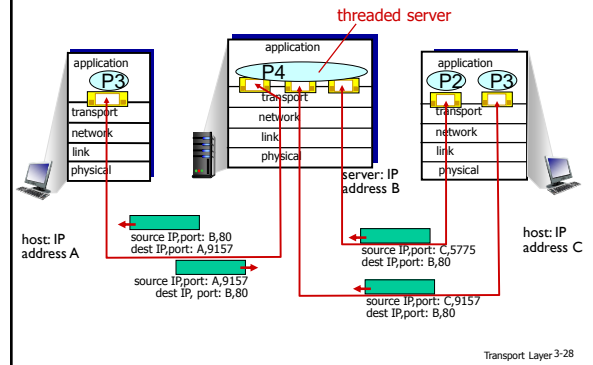
- ❖ TCP socket identified by 4-tuple:
 - source IP address
 - source port number
 - dest IP address
 - dest port number
- ❖ demux: receiver uses all four values to direct segment to appropriate socket
- ❖ server host may support many simultaneous TCP sockets:
 - each socket identified by its own 4-tuple
- ❖ web servers have different sockets for each connecting client
 - non-persistent HTTP will have different socket for each request

Transport Layer 3-26

Connection-oriented demux: example



Connection-oriented demux: example

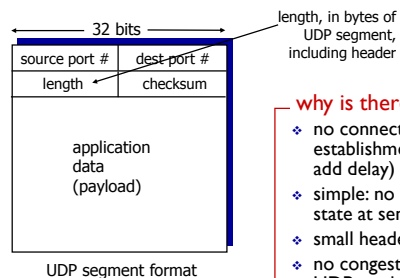


UDP: User Datagram Protocol [RFC 768]

- ❖ “no frills,” “bare bones” Internet transport protocol
- ❖ “best effort” service, UDP segments may be:
 - lost
 - delivered out-of-order to app
- ❖ **connectionless:**
 - no handshaking between UDP sender, receiver
 - each UDP segment handled independently of others
- ❖ UDP use:
 - streaming multimedia apps (loss tolerant, rate sensitive)
 - DNS
 - SNMP
- ❖ reliable transfer over UDP:
 - add reliability at application layer
 - application-specific error recovery!

Transport Layer 3-29

UDP: segment header



why is there a UDP?

- ❖ no connection establishment (which can add delay)
- ❖ simple: no connection state at sender, receiver
- ❖ small header size
- ❖ no congestion control: UDP can blast away as fast as desired

Transport Layer 3-30

UDP checksum

Goal: detect “errors” (e.g., flipped bits) in transmitted segment

sender:

- ❖ treat segment contents, including header fields, as sequence of 16-bit integers
- ❖ checksum: addition (one’s complement sum) of segment contents
- ❖ sender puts checksum value into UDP checksum field

receiver:

- ❖ compute checksum of received segment
- ❖ check if computed checksum equals checksum field value:
 - NO - error detected
 - YES - no error detected. *But maybe errors nonetheless? More later in Sec 5.2*

Transport Layer 3-31

Checksum: example

example: add two 16-bit integers

```


      1 1 1 0 0 1 1 0 0 1 1 0 0 1 1 0
      1 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1
      -----
wraparound ① 1 0 1 1 1 1 0 1 1 1 0 1 1 1 0 1 1
      sum      1 0 1 1 1 1 0 1 1 1 0 1 1 1 0 0
      checksum 0 1 0 0 0 1 0 0 0 1 0 0 0 0 1 1
    
```

Note: when adding numbers, a carryout from the most significant bit needs to be added to the result

Transport Layer 3-32

Principles of reliable data transfer

- ❖ important in application, transport, link layers
 - top-10 list of important networking topics!




(a) provided service

- ❖ characteristics of unreliable channel will determine complexity of reliable data transfer protocol (rdt)

Transport Layer 3-33

Principles of reliable data transfer

- ❖ important in application, transport, link layers
 - top-10 list of important networking topics!




(a) provided service (b) service implementation

- ❖ characteristics of unreliable channel will determine complexity of reliable data transfer protocol (rdt)

Transport Layer 3-34

Principles of reliable data transfer

- ❖ important in application, transport, link layers
 - top-10 list of important networking topics!



(a) provided service (b) service implementation


- ❖ characteristics of unreliable channel will determine complexity of reliable data transfer protocol (rdt)

Transport Layer 3-35

Reliable data transfer: getting started

`rdt_send()`: called from above, (e.g., by app.). Passed data to deliver to receiver upper layer

`deliver_data()`: called by `rdt` to deliver data to upper



`udt_send()`: called by `rdt`, to transfer packet over unreliable channel to receiver


`rdt_rcv()`: called when packet arrives on rcv-side of channel

Transport Layer 3-36

Reliable data transfer: getting started

we'll:


- incrementally develop sender, receiver sides of reliable data transfer protocol (rdt)
- consider only unidirectional data transfer
 - but control info will flow on both directions!
- use finite state machines (FSM) to specify sender, receiver



Transport Layer 3-37

rdt1.0: reliable transfer over a reliable channel

- underlying channel perfectly reliable
 - no bit errors
 - no loss of packets
- separate FSMs for sender, receiver:
 - sender sends data into underlying channel
 - receiver reads data from underlying channel



Transport Layer 3-38

rdt2.0: over a channel with bit errors

- underlying channel may flip bits in packet
 - checksum to detect bit errors
- the question: how to recover from errors:

How do humans recover from "errors" during conversation?

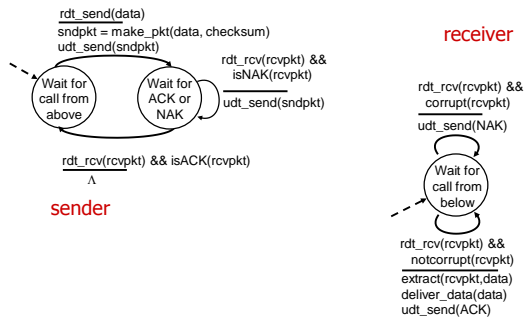
Transport Layer 3-39

rdt2.0: channel with bit errors

- underlying channel may flip bits in packet
 - checksum to detect bit errors
- the question: how to recover from errors:
 - acknowledgements (ACKs):** receiver explicitly tells sender that pkt received OK
 - negative acknowledgements (NAKs):** receiver explicitly tells sender that pkt had errors
 - sender retransmits pkt on receipt of NAK
- new mechanisms in rdt2.0 (beyond rdt1.0):
 - error detection
 - feedback: control msgs (ACK,NAK) from receiver to sender

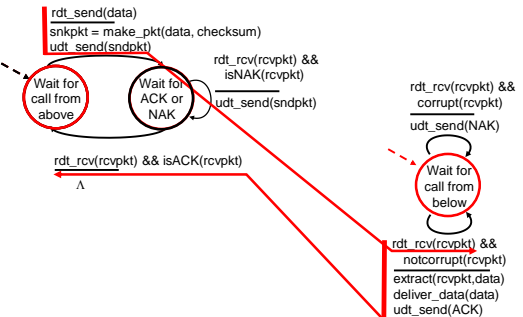
Transport Layer 3-40

rdt2.0: FSM specification



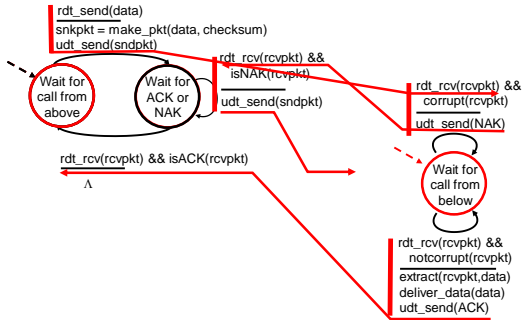
Transport Layer 3-41

rdt2.0: operation with no errors



Transport Layer 3-42

rdt2.0: error scenario



Transport Layer 3-43

rdt2.0 has a fatal flaw!

what happens if ACK/NAK corrupted?

- ❖ sender doesn't know what happened at receiver!
- ❖ can't just retransmit: possible duplicate

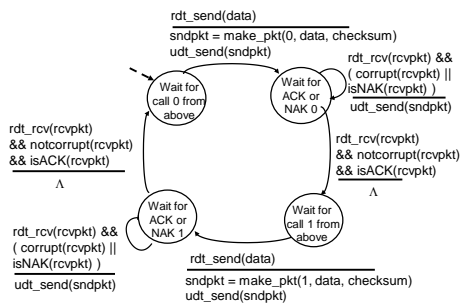
handling duplicates:

- ❖ sender retransmits current pkt if ACK/NAK corrupted
- ❖ sender adds *sequence number* to each pkt
- ❖ receiver discards (doesn't deliver up) duplicate pkt

stop and wait
sender sends one packet,
then waits for receiver
response

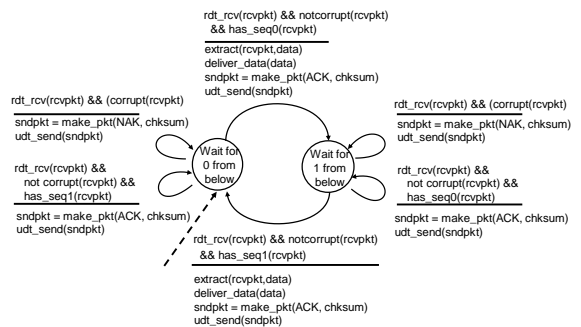
Transport Layer 3-44

rdt2.1: sender, handles garbled ACK/NAKs



Transport Layer 3-45

rdt2.1: receiver



Transport Layer 3-46

rdt2.1: summary

sender:

- ❖ seq # added to pkt
- ❖ two seq. #'s (0,1) will suffice. Why?
- ❖ must check if received ACK/NAK corrupted
- ❖ twice as many states
 - state must "remember" whether "expected" pkt should have seq # of 0 or 1

receiver:

- ❖ must check if received packet is duplicate
 - state indicates whether 0 or 1 is expected pkt seq #
- ❖ note: receiver can *not* know if its last ACK/NAK received OK at sender

Transport Layer 3-47

Next

- ❖ Midterm on Chapters 1 and 2
 - Exclude Sec 1.7, 2.4, 2.7, and 2.8
- ❖ ~10% concepts addressed in class
- ❖ ~90% problems requiring math models
 - In addition to Assignments 1 and 2, Tutorials 1 to 4, reading from the book, make sure you have no doubts on the following problems:
 - Ch1: P2-P17, P19-P30, and P32
 - Ch2: P1, P7-P11, P20-P33
 - Note that our reference is the 5th edition

2-48