

Reliable Data Transfer

Kai Shen

10/14/2009

CSC 257/457 - Fall 2009

1

Reliable Data Transfer

- What is **reliable data transfer**?
 - guaranteed arrival
 - no error
 - in order delivery
- Why is it difficult?
 - unreliable underlying communication channel, which can be **lossy, error-prone, and deliver packets out of order**
- Where is it used in computer networks?
 - reliable transport service on top of unreliable network layer
 - reliable data link service on top of unreliable physical layer

10/14/2009

CSC 257/457 - Fall 2009

2

Principles of Reliable Data Transfer

- Characteristics of unreliable channel will determine complexity of reliable data transfer protocol
 - e.g., delay in the channel is bounded in physical layer, not so for network layer
- Other services may interact with RDT protocol
 - e.g., flow control, congestion control
- Here we study widely applicable RDT principles
 - we don't make assumptions about the unreliable channel
 - we don't consider interaction with other services
- Later we see what RDT is like in practice
 - in a transport layer protocol - TCP

10/14/2009

CSC 257/457 - Fall 2009

3

Outline

- Overview of reliable data transfer
- a **correct** protocol: stop-and-wait
 - one packet at a time
- an **efficient** protocol: sliding window
 - multiple packets simultaneously

10/14/2009

CSC 257/457 - Fall 2009

4

Deal with Errors

- First deal with errors, later deal with packet loss.
- ACK-based solution: receiver check errors
 - if correct, send back positive ACK
 - otherwise, send back negative NAK

What if ACK or NAK is corrupted?

10/14/2009 CSC 257/457 - Fall 2009 5

What if ACK or NAK is corrupted?

- Solution 1: creating special acknowledgments for ACKs/NAKs. What if they get corrupted too??
- Solution 2: treat corrupted acknowledgements as NAKs. **Duplicated packets!!**
- To solve duplicated packets: **sequence number** for each packet.

How many sequence numbers do we need here?

10/14/2009 CSC 257/457 - Fall 2009 6

Deal with Packet Loss: Timeouts

- Early timeout => duplicated packet => sequence number. (not likely for data link protocol)

10/14/2009 CSC 257/457 - Fall 2009 7

Deal with Duplicated ACKs

- Solution: each ACK carries sequence number.
- With timeout, NAK is not necessary any more.

10/14/2009 CSC 257/457 - Fall 2009 8

Stop-and-Wait

Now we have a *correct* protocol:

- Allow one outstanding (un-ACKed) packet - *stop-and-wait*
- By the way, we haven't talked about in-order delivery.

10/14/2009 CSC 257/457 - Fall 2009 9

Efficiency of Stop-and-Wait

Example:

- Packet size $L = 1\text{KB}$ (8kbits),
- Transmission speed $R = 1\text{Gbps}$,
- Roundtrip prop. delay $\text{RTT} = 30\text{ms}$.

\Rightarrow **0.027% efficiency!**

$$\text{efficiency} = \frac{L / R}{\text{RTT} + L / R}$$

10/14/2009 CSC 257/457 - Fall 2009 10

Pipelined Protocols

Pipelining: sender allows multiple, "in-flight", yet-to-be-acknowledged packets

(a) a stop-and-wait protocol in operation (b) a pipelined protocol in operation

10/14/2009 CSC 257/457 - Fall 2009 11

Pipelining: Increased Efficiency

Increase utilization by a factor of 3!

$$\text{Channel util efficiency} = \frac{3 * L / R}{\text{RTT} + L / R} = \frac{.000024}{.030008} = 0.08\%$$

10/14/2009 CSC 257/457 - Fall 2009 12

A pipelined protocol: Sliding Window

- Allow multiple outstanding (un-ACKed) packets
- Upper bound on un-ACKed packets, called **window**

Two variations: go-back-N, and selective repeat.

10/14/2009 CSC 257/457 - Fall 2009 13

Go-Back-N: Sender

- "window" of up to N consecutive un-ACKed packets allowed; limit send buffer space

- cumulative ACK - ACK with seq #n stands for ACKs all packets up to, including seq #n
- recv ACKs in send window
 - sliding send window
- timer for each in-flight packet (or first in-flight packet)
- packet with seq #n timeouts:
 - retransmit #n and all higher seq # packets in window (**buffering**)

10/14/2009 CSC 257/457 - Fall 2009 14

Go-Back-N: Receiver

- In-order packet:
 - Acknowledge it.
- Out-of-order packet:
 - Discard!
 - Optional: Re-ACK packet with highest in-order seq # (sort of a NACK)
 - alert sender something is wrong through duplicated ACKs
 - not critical for protocol correctness; but may improve performance

10/14/2009 CSC 257/457 - Fall 2009 15

GBN in Action

Drawback:

- Resend out-of-order packets

To fix it:

- Receiver buffering
- Selective acknowledgement

10/14/2009 CSC 257/457 - Fall 2009 16

Selective Repeat

- receiver
 - buffers out-of-order packets for eventual in-order delivery to upper layer
 - individually acknowledges all correctly received packets
- sender
 - maintains timer for each un-ACKed packet
 - only resends packets whose timers expire before ACKs are received

10/14/2009 CSC 257/457 - Fall 2009 17

Selective Repeat: Sender, Receiver Windows

10/14/2009 CSC 257/457 - Fall 2009 18

Selective Repeat in Action

19

Selective Repeat Implementation

sender

data from above:

- if there is available slot in window, send pkt

timeout(n):

- resend pkt n, restart timer

ACK(n):

- mark pkt n as received
- if n is smallest unACKed pkt, advance window base to next unACKed seq # (sliding!)

receiver

pkt n in [rcvbase, rcvbase+N-1]

- send ACK(n)
- in-order: deliver (also deliver buffered, in-order pkts, advance window to next not-yet-received pkt (sliding!))
- out-of-order: buffer


pkt n in [rcvbase-N, rcvbase-1]

- ACK(n)

otherwise:

- ignore


10/14/2009 CSC 257/457 - Fall 2009 20



Sequence Number Space

- **SeqNum** field is finite; sequence numbers wrap around
 - 0, 1, ..., S-1, 0, 1, ..., S-1, 0, 1, ...
- For a given window size N, how large a sequence number space is enough?
 - out-of-order channel: an old packet carrying the same sequence number confused with a packet of current concern

10/14/2009 CSC 257/457 - Fall 2009 21



Disclaimer

- Parts of the lecture slides contain original work of James Kurose, Larry Peterson, and Keith Ross. The slides are intended for the sole purpose of instruction of computer networks at the University of Rochester. All copyrighted materials belong to their original owner(s).

10/14/2009 CSC 257/457 - Fall 2009 22