

CSE/EE 461 – Lecture 13

Sliding Windows and ARQ

David Wetherall
djw@cs.washington.edu

Last Time

- We finished up the Network layer
 - Internetworks (IP)
 - Routing (DV/RIP, LS/OSPF)
 - Scalable addressing/routing (BGP, CIDR)
- It was all about routing: how to provide end-to-end delivery of packets.

Application
Presentation
Session
Transport
Network
Data Link
Physical

This Time

- We begin on the Transport layer
- Focus
 - How do we send information reliably?
- Topics
 - The Transport layer
 - Acknowledgements and retransmissions (ARQ)
 - Sliding windows

Application
Presentation
Session
Transport
Network
Data Link
Physical

The Transport Layer

- Builds on the services of the Network layer
- Communication between processes running on hosts
 - Naming/Addressing
- Stronger guarantees of message delivery
 - Reliability

Example – Common Properties

TCP

- Guaranteed delivery
- In-order delivery
- Single delivery
- Arbitrarily long messages
- Synchronization
- Flow control
- Multiple processes


IP

- Lost packets
- Reordered packets
- Duplicate packets
- Limited size packets

Internet Transport Protocols

- UDP
 - Datagram abstraction between processes
 - With error detection
- TCP
 - Bytestream abstraction between processes
 - With reliability
 - Plus congestion control (later!)

Automatic Repeat Request (ARQ)

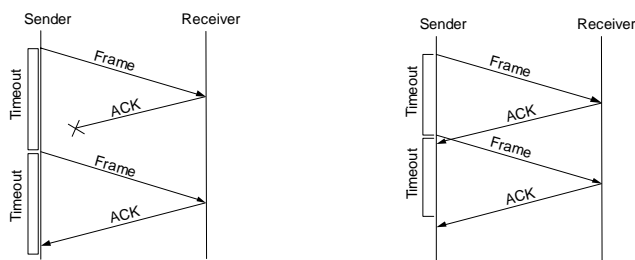


- Packets can be corrupted or lost. How do we add reliability?
- Acknowledgments (ACKs) and retransmissions after a timeout
- ARQ is generic name for protocols based on this strategy

djw // CSE/EE 461, Autumn 2002

L13.7

The Need for Sequence Numbers




- In the case of ACK loss (or poor choice of timeout) the receiver can't distinguish this message from the next
 - Need to understand how many packets can be outstanding and number the packets; here, a single bit will do

djw // CSE/EE 461, Autumn 2002

L13.8

Stop-and-Wait

- Only one outstanding packet at a time
- Also called alternating bit protocol



djw // CSE/EE 461, Autumn 2002

L13.9

Limitation of Stop-and-Wait




- Lousy performance if wire time \ll prop. delay
 - How bad? You do the math
- Want to utilize all available bandwidth
 - Need to keep more data "in flight"
 - How much? Remember the bandwidth-delay product?
- Leads to Sliding Window Protocol

djw // CSE/EE 461, Autumn 2002


L13.10

Sliding Window – Sender




- Window bounds outstanding data
 - Implies need for buffering at sender
- “Last” ACK applies to in-order data
- Sender maintains timers too
 - Go-Back-N: one timer, send all unacknowledged on timeout
 - Selective Repeat: timer per packet, resend as needed

Sliding Window – Timeline



Sliding Window – Receiver



- Receiver buffers too:
 - data may arrive out-of-order
 - or faster than can be consumed (flow control)
- Receiver ACK choices:
 - Individual, Cumulative (TCP), Selective (newer TCP), Negative

Sliding Window Functions

- Sliding window is a mechanism
- It supports multiple functions:
 - Reliable delivery
 - In-order delivery
 - Flow control


Flow Control

- Sender must transmit data no faster than it can be consumed by the receiver
 - Receiver might be a slow machine
 - App might consume data slowly
- Implement by adjusting the size of the sliding window used at the sender based on receiver feedback about available buffer space
 - This is the purpose of the Advertised Window field

djw // CSE/EE 461, Autumn 2002

L13.15


Sender and Receiver Buffering



djw // CSE/EE 461, Autumn 2002

L13.16


Example – Exchange of Packets



djw // CSE/EE 461, Autumn 2002

L13.17

Example – Buffer at Sender



djw // CSE/EE 461, Autumn 2002

L13.18

Key Concepts

- Transport layer allows processes to communicate with stronger guarantees, e.g., reliability
- Basic reliability is provided by ARQ mechanisms
 - Stop-and-Wait through Sliding Window plus retransmissions