CSCI-1680 Transport Layer I

Nick DeMarinis

Based partly on lecture notes by Rodrigo Fonseca, David Mazières, Phil Levis, John Jannotti

#### Administrivia

- Later today: Look for message about IP grading
  - Meeting slots first week after break (and during break)
- TCP: Draft of assignment out today
  - Read it over before break, start when we get back
- Summer/UTA hiring: Expect a message from me today/ tomorrow

Today

Light overview of the transport layer and TCP

- Why we need TCP
- What components are involved
- What you will do in the project

#### Transport Layer



- Transport protocols sit on top of network layer
- Problem solved: communication among processes
  - Application-level multiplexing ("ports")
  - Error detection, reliability, etc.











## Basic transport: UDP



#### User Datagram Protocol

- Unreliable datagram service
- Adds multiplexing (via ports) and nothing else
- Checksum is pretty useless

## Next Problem: Reliability

ACK

We talked briefly about link-layer reliability:

	Problem		Mechan	ism	
$\square$	Dropped Packets	Ac	knowledgments + <sup>-</sup>	Timeout	
ſ	Duplicate Packets	Se	quence Numbers		
1	Packets out of order	Re	ceiver Window		
	Maximizing throughput	Slie	ling Window (Pipelining)		
		Ľ	71		1

## Next Problem: Reliability

#### We talked briefly about link-layer reliability:

Problem	Mechanism
Dropped Packets	Acknowledgments + Timeout
Duplicate Packets	Sequence Numbers
Packets out of order	Receiver Window
Maximizing throughput	Sliding Window (Pipelining)

• Single link: things were easy... 🙂

## Transport Layer Reliability

- Extra difficulties
  - Multiple hosts
  - Multiple hops
  - Multiple potential paths

## Transport Layer Reliability

- Extra difficulties
  - Multiple hosts
  - Multiple hops
  - Multiple potential paths
- What does this mean?
  - Multiple opportunities for failure
  - Hosts have different resources
  - Varying RTTs

## Extra Difficulties (cont.)

- Out of order packets
  - Not only because of drops/retransmissions
  - Can get very old packets (up to 120s), must not get confused

## Extra Difficulties (cont.)

- Out of order packets
  - Not only because of drops/retransmissions
  - Can get very old packets (up to 120s), must not get confused
- Unknown resources at other end
  - Must be able to discover receiver buffer: flow control

STOP/LIMIT DATA

## Extra Difficulties (cont.)

• Out of order packets



- Not only because of drops/retransmissions
- Can get very old packets (up to 120s), must not get confused
- Unknown resources at other end
  - Must be able to discover receiver buffer: flow control
- Unknown resources in the network
  - Should not overload the network
  - But should use as much as safely possible to maximize throughput

"FAIRNESS" IN USE OF CARCE RW

74 MOVIDES TO NOST ON RECV. SEND DATA 70 TP (KENLUEC) (YOUR NODE) TCP STACK TCP STACK R. Buis. BUFFER SEND SENDS SELMENTS SEGMENTS SEND ACK (1) <del>ا ب</del>ا 42 浙





 Service model: "reliable, connection oriented, full duplex ordered byte stream"



- Service model: "reliable, connection oriented, full duplex ordered byte stream"
- Flow control: If one end stops reading, writes at other eventually stop/fail



- Service model: "reliable, connection oriented, full duplex ordered byte stream"
- Flow control: If one end stops reading, writes at other eventually stop/fail
- Congestion control: Keeps sender from overloading the network

#### TCP

- Specification
  - RFC 793 (1981), RFC 1222 (1989, some corrections), RFC 5681 (2009, congestion control), ...
- Was born coupled with IP, later factored out - NOPS IN THE MIDDLE JUST FORWARD BACKETS NOT AWARE OF TCP.
- End-to-end protocol – Minimal assumptions on the network - All mechanisms run on the end points
- What if you had link-layer reliability instead?

# Why not provide X on the network layer?

X = Reliability, security, message ordering...

- Cost
  - These functionalities are not free: don't burden those who don't need them

# Why not provide X on the network layer?

X = Reliability, security, message ordering...

- Cost
  - These functionalities are not free: don't burden those who don't need them
- Conflicting

- Timeliness and in-order delivery, for example

# Why not provide X on the network layer?

X = Reliability, security, message ordering...

- Cost
  - These functionalities are not free: don't burden those who don't need them
- Conflicting
  - Timeliness and in-order delivery, for example
- Insufficient
  - Example: reliability

#### End-to-end argument

- Functions placed at lower levels of a system may be redundant or of little value
  - They may **need** to be performed at a higher layer anyway
- But they may be justified for performance reasons

## End-to-end argument

- Functions placed at lower levels of a system may be redundant or of little value
  - They may **need** to be performed at a higher layer anyway
- But they may be justified for performance reasons
  - Or just because they provide most of what is needed
  - Example: retransmissions
- Takeaway: weigh the costs and benefits at each layer

#### TCP Header



#### Header Fields

- Ports: multiplexing
- Sequence number
  - Correspond to bytes, not packets!
- Acknowledgment Number
  - Next expected sequence number
- Window: willing to receive
  - Lets receiver limit SWS (even to 0) for flow control
- Data Offset: # of 4 byte (header + option bytes)
- Flags, Checksum, Urgent Pointer

# Header Flags

- URG: whether there is urgent data
- ACK: ack no. valid (all but first segment)
- PSH: push data to the application immediately
- RST: reset connection
- SYN: synchronize, establishes connection
- FIN: close connection



# Establishing a Connection



- Three-way handshake
  - Two sides agree on respective initial sequence nums
- If no one is listening on port: server sends RST
- If server is overloaded: ignore SYN
- If no SYN-ACK: retry, timeout

#### **Connection Termination**

- FIN bit says no more data to send
  - Caused by close or shutdown
  - Both sides must send FIN to close a connection
- Typical close FIN\_WAIT\_1 FIN\_WAIT\_2 FIN\_WAIT\_2 FIN\_WAIT\_2 FIN ACK CLOSE\_WAIT Close LAST\_ACK CLOSED CLOSED









#### Next class

• Sending data over TCP