CS640: Introduction to Computer Networks

Aditya Akella

Lecture 15

TCP – II -

Connection Set-up and Congestion Control

TCP Packet

Reliable, In-order, Connection oriented, Byte stream abstraction

Flags from MSB to LSB: URG
ACK
PSH
RST
SYN
FIN

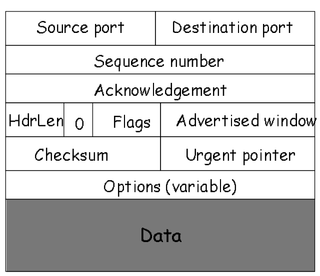

2

Sequence and Acknowledge Numbers

- Sequence number → byte num of first byte in payload
- Acknowledgement number
  - TCP is full duplex
  - Sequence number of next byte expected in reverse direction

3

Advertised Window

- Used for "flow control"
  - Different from "congestion control", which we will see in second half of today's lecture
- Both sender and receiver advertise window
  - Sender action:
    lastSent – lastACK <= Receiver's advertised window

4

Establishing Connection: Three-Way handshake

- Each side notifies other of starting sequence number it will use for sending
  - Why not simply chose 0?
    - Must avoid overlap with earlier incarnation
    - Security issues
- Each side acknowledges other's sequence number
  - SYN-ACK: Acknowledge sequence number + 1
- Can combine second SYN with first ACK

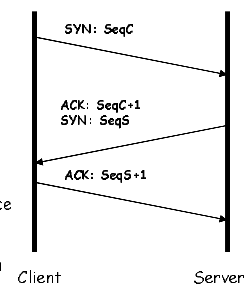

5

Tearing Down Connection

- Either side can initiate tear down
  - Send FIN signal
  - "I'm not going to send any more data"
- Other side can continue sending data
  - Half open connection
  - Must continue to acknowledge
- Acknowledging FIN
  - Acknowledge last sequence number + 1

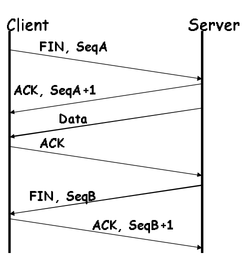

6

TCP State Diagram: Connection Setup

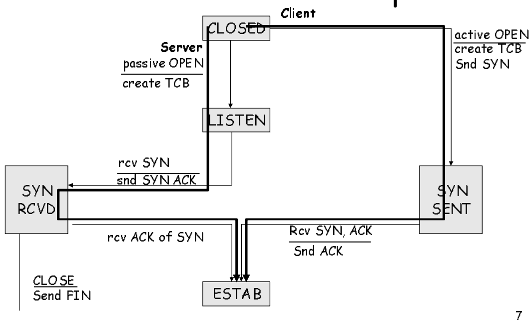

7

State Diagram: Connection Tear-down

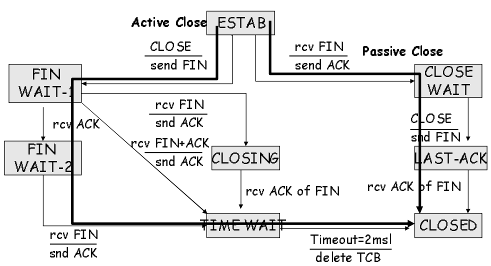

8

Congestion

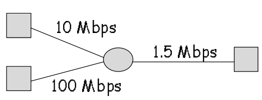

- Different sources compete for resources inside network
- Why is it a problem?
  - Sources are unaware of current state of resource
  - Sources are unaware of each other
- Manifestations:
  - Lost packets (buffer overflow at routers)
  - Long delays (queuing in router buffers)
  - Can result in effective throughput less than bottleneck link (1.5Mbps for the above topology) → a.k.a. congestion collapse

9

Causes & Costs of Congestion

- Four senders – multihop paths
- Timeout/retransmit

<u>Q:</u> What happens as rate increases?

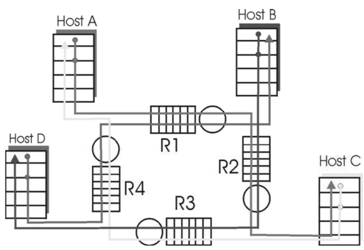

10

Causes & Costs of Congestion

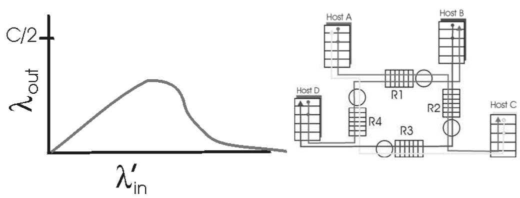

- When packet dropped, any upstream transmission capacity used for that packet was wasted!

11

Congestion "Collapse"

- Definition: *Unchecked Increase in network load results in decrease of useful work done*
  - Fewer and fewer useful packets carried in network
- Many possible causes
  - Spurious retransmissions of packets still in flight
    - Classical congestion collapse
  - Undelivered packets
    - Packets consume resources and are dropped elsewhere in network

12

Congestion Control and Avoidance

- A mechanism which:
  - Uses network resources efficiently
  - Preserves fair network resource allocation
  - Controls or Avoids congestion

13

Approaches Towards Congestion Control

- Two broad approaches towards congestion control:

- End-end congestion control:
  - No explicit feedback from network
  - Congestion inferred from end-system observed loss, delay
  - Approach taken by TCP
  - Problem: approximate, possibly inaccurate

- Network-assisted congestion control:
  - Routers provide feedback to end systems
    - Single bit indicating congestion (SNA, DECbit, TCP/IP ECN, ATM)
    - Explicit rate sender should send at
  - Problem: makes routers complicated

14

End-End Congestion Control

- So far: TCP sender limited by available buffer size at receiver
  - Receiver flow control
  - "receive window" or "advertised window"

- To accommodate network constraints, sender maintains a "congestion window"
  - Reflects *dynamic* state of the network
  - ***Max outstanding packets ≤ min {congestion window, advertised window}***

- When receiver window is very large, congestion window determines how fast sender can send
  - ***Speed = CWND/RTT (roughly)***

15

TCP Congestion Control

- Very simple mechanisms in network
  - FIFO scheduling with shared buffer pool
  - Feedback through packet drops
- End-host TCP interprets drops as signs of congestion and slows down → reduces size of congestion window
- But then, periodically probes – or increases congestion window
  - To check whether more bandwidth has become available

16

Congestion Control Objectives

- Simple router behavior
- Distributed-ness
- Efficiency: $\Sigma x_i(t)$ close to system capacity
- Fairness: equal (or propotional) allocation
  - Metric = $(\Sigma x_i)^2/n(\Sigma x_i^2)$
- Convergence: control system must be stable

17

Linear Control

- Many different possibilities for reaction to congestion and probing
  - Examine simple linear controls
    - Window(t + 1) = a + b Window(t)
    - Different a_i/b_i for increase and a_d/b_d for decrease
- Various reaction to signals possible
  - Increase/decrease additively
  - Increased/decrease multiplicatively
  - Which of the four combinations is optimal?
    - Consider two end hosts vying for network bandwidth

18

Additive Increase/Decrease

- Both X_1 and X_2 increase/ decrease by the same amount over time
  - Additive increase improves fairness and additive decrease reduces fairness

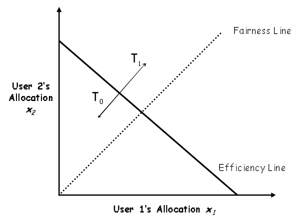

19

Multiplicative Increase/Decrease

- Both X_1 and X_2 increase by the same factor over time
  - Extension from origin – constant fairness

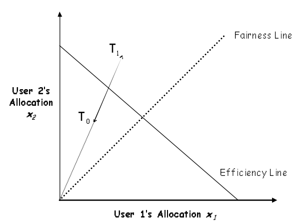

20

Convergence to Efficiency

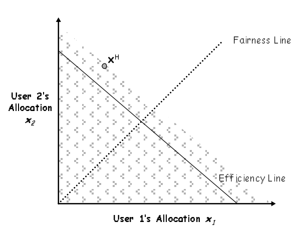

21

Distributed Convergence to Efficiency

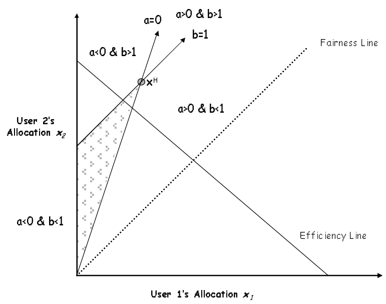

22

Convergence to Fairness

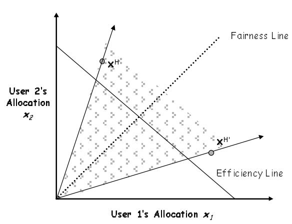

23

Convergence to Efficiency & Fairness

- Intersection of valid regions
- For decrease: a=0 & b < 1

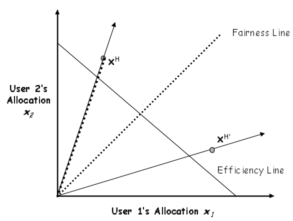

24

What is the Right Choice?

- Constraints limit us to AIMD
  - Can have multiplicative term in increase (MAIMD)
  - AIMD moves towards optimal point

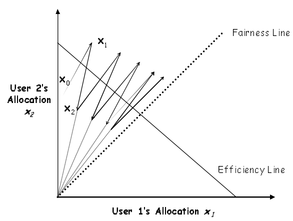

25

Summary

- Significance of fields in TCP packet
  - TCP is full duplex
  - Data can go in either direction
- TCP connection set-up – three-way hand-shake
  - Also, teardown
- Costs of congestion
  - Delay, loss, useless work…
  - Cure: congestion control
- TCP uses AIMD congestion control

26