CS640: Introduction to Computer Networks

Aditya Akella
Lecture 15
TCP Congestion Control

TCP State Diagram: Connection Setup

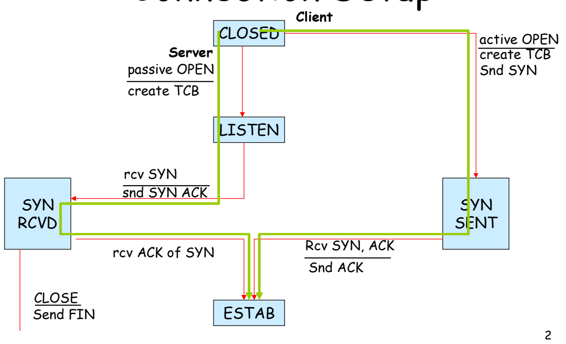

2

State Diagram: Connection Tear-down

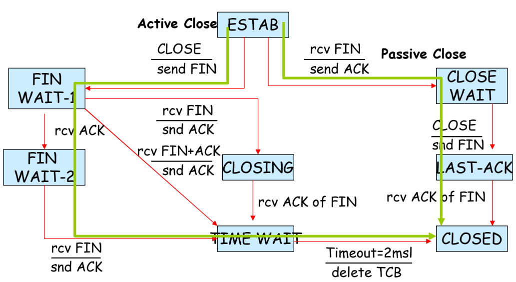

Time_Wait state is necessary in case the final ack was lost.

3

From the Previous Lecture: TCP Persist in Sliding Window Flow Control

- What happens if window is 0?
  - Receiver updates window when application reads data
  - What if this update is lost?
- TCP Persist state
  - Sender periodically sends 1 byte packets
  - Receiver responds with ACK even if it can't store the packet

4

Congestion

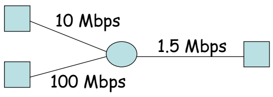

- Different sources compete for resources inside network
- Why is it a problem?
  - Sources are unaware of current state of resource
  - Sources are unaware of each other
- Manifestations:
  - Lost packets (buffer overflow at routers)
  - Long delays (queuing in router buffers)
  - Can result in effective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

5

Causes & Costs of Congestion

- Four senders – multihop paths
- Timeout/retransmit

Q: What happens as rate increases?

Only output buffers used

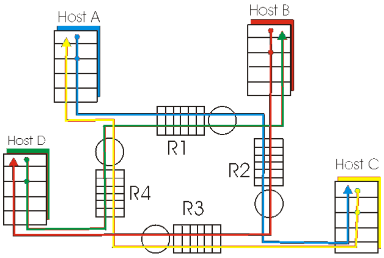

6

Causes & Costs of Congestion

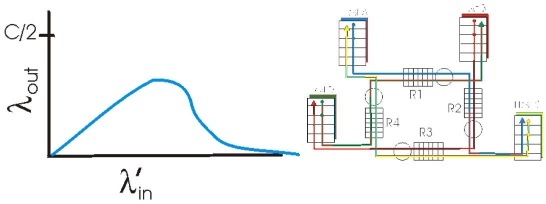

- When packet dropped, any upstream transmission capacity used for that packet was wasted!

7

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

8

Congestion Control and Avoidance

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

9

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

10

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)***

11

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

12

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

13

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

14

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

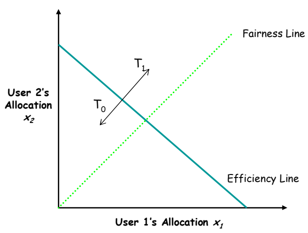

15

Multiplicative Increase/Decrease

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

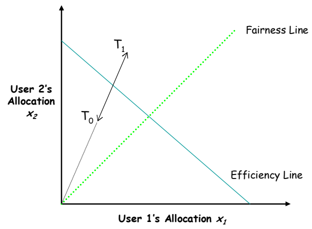

16

Convergence to Efficiency

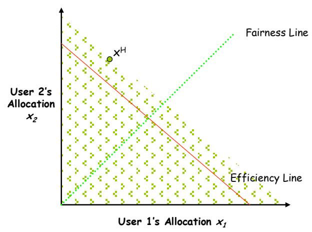

17

Distributed Convergence to Efficiency

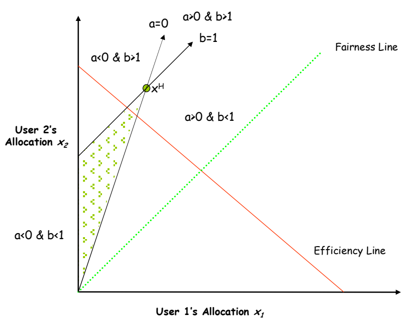

18

Convergence to Fairness

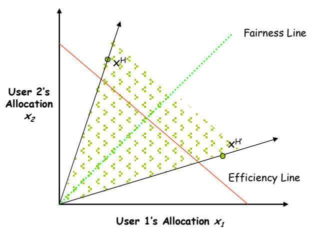

19

Convergence to Efficiency & Fairness

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

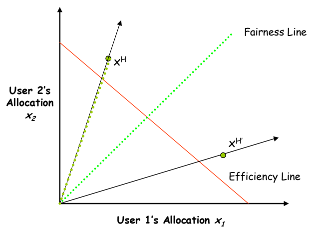

20

What is the Right Choice?

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

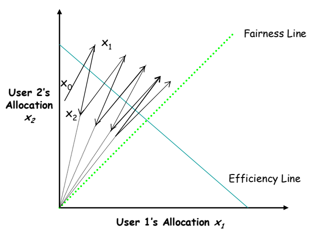

21