Announcements

- Midterms
  - Mt #1 Tuesday March 6
  - Mt #2 Tuesday April 15
  - Final project design due **April 11**
- Midterm #1
  - Chapters 1 & 2
  - Chapter 5 (to 5.2)

Congestion

- Too much traffic can destroy performance
  - goal is to permit the network to operate near link capacity
  - can reach a knee in the packets sent vs. delivered curve
- Sources
  - all traffic is destined for a single out link
    - backup in traffic consumes buffers
    - other (cross traffic) will not get through due to lack of buffers
  - slow router CPU
    - can’t service all requests at link speed
      - links still backup
- Often feeds on itself
  - queuing delays can cause packets to timeout
    - introduces more traffic due to re-transmissions

Congestion Control

- Two possible approaches
  - open loop: prevent congestion from every happening
    - tends to be conservative and result in under utilizaion
  - closed loop: detect and correct
    - some congestion will still occur until it is corrected
- Open loop
  - request resources before using them
  - global (or regional) resource allocation
    - responds yes or no to each request for service
- Closed loop
  - monitor network to detect congestion
  - pass information back to location where action can be taken
  - adjust system operation to correct the problem

Responding to Congestion

- Add more resources
  - dialup network: start making additional connections
  - SMDS: request additional bandwidth from provider
  - split traffic: use all routes not just optimal
- Decrease load
  - deny service to some users: based on priorities
  - degrade service to some or all users
  - require users to schedule their traffic

Traffic Shaping

- Traffic tends to be bursty
  - great variation between min and max bandwidth used
  - this uncertainty leads to inefficient use of the network
- Flow Specification
  - user proposes a specific probability distribution
    - maximum packet size
    - transmission rate (min, max, or mean)
    - maximum delay
    - maximum delay variation (jitter)
    - quality guarantee (how strong is this agreement)
  - network can
    - agree to request
    - refuse it
    - counter offer

Leaky Bucket

- buffer accepts traffic at link rate
  - buffer has a bounded size (limits burst size that is accepted)
- output is limited to a lower rate
  - traffic is constrained to this rate

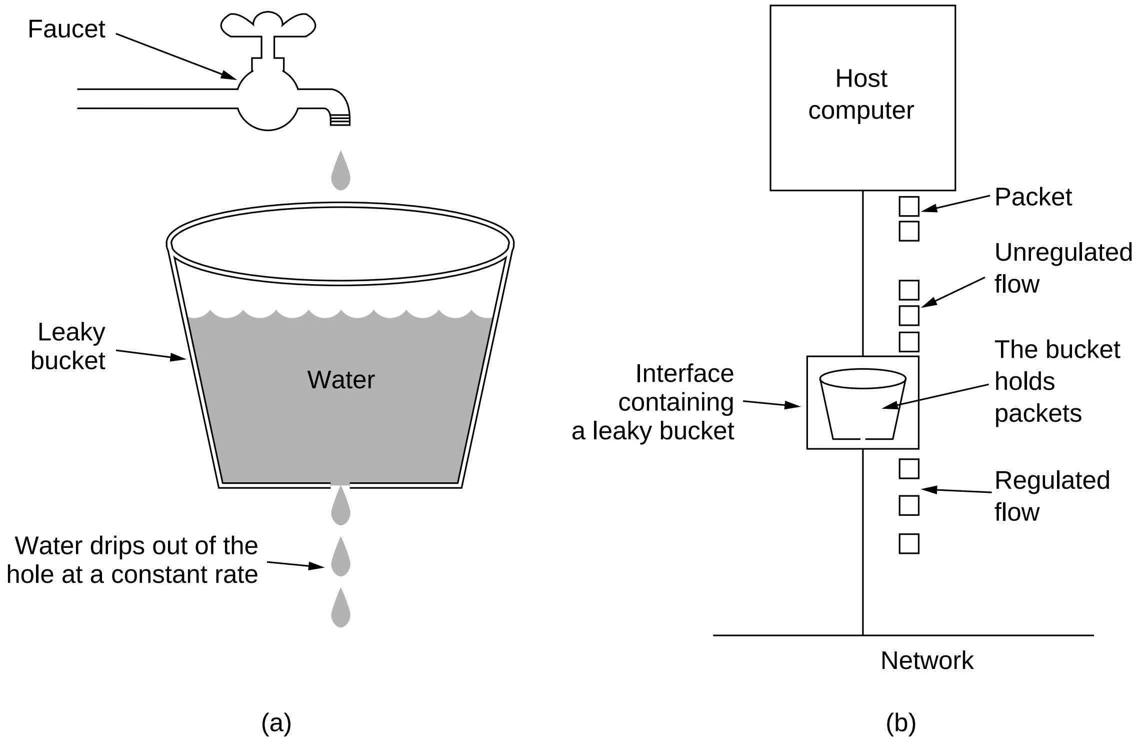

From: *Computer Networks*, 3rd Ed. by Andrew S. Tanenbaum, (c)1996 Prentice Hall.

Token Bucket

- Bucket hold tokens (generated one every T seconds)
- Can save up to a fixed limit of n tokens
- When traffic arrives, it must a have token to be sent

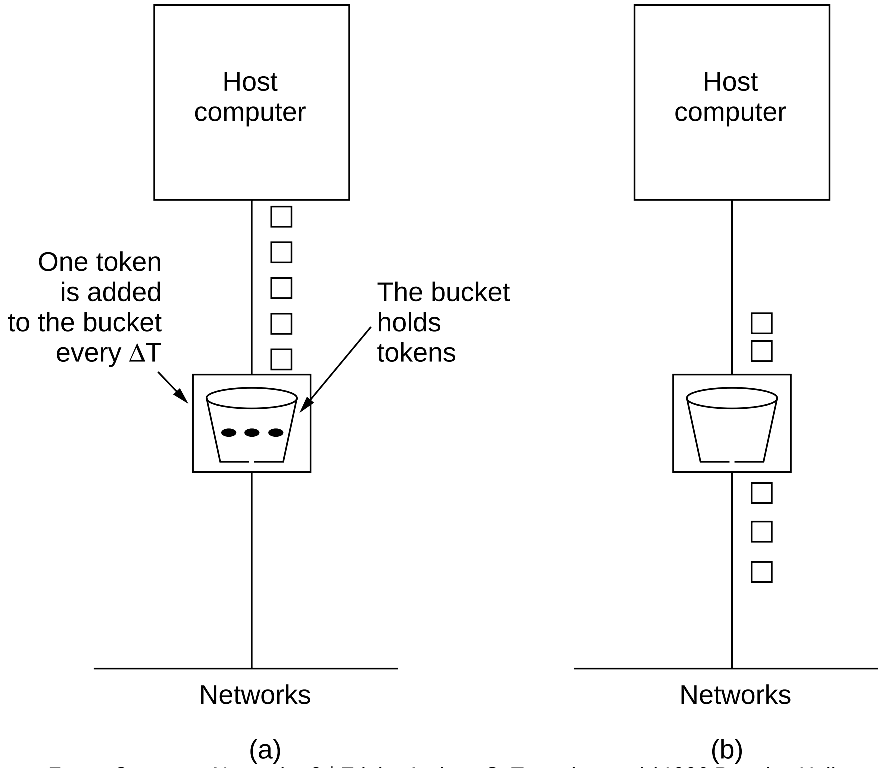

From: *Computer Networks*, 3rd Ed. by Andrew S. Tanenbaum, (c)1996 Prentice Hall.

- Max burst rate

  C - capacity of bucket

  S - burst length in seconds

  M - max output rate

  p - token credit rate

  $C + pS = MS$

Congestion Control with Virtual Circuits

- Admission control
  - once traffic reaches a threshold, don’t admit more VCs
  - doesn’t correct current problem, but prevents additional congestion
- Alter routes
  - admit new connections
  - route them around “trouble” areas
- Negotiate traffic
  - establish parameters for volume and shape of traffic

Choke Packets

- Monitor link utilization
  - keep an estimate (*u*) of average utilization over time
  - $u_{new} = au_{old} + (1 - a)f$
    - f is a 0/1 sampling of link state
    - a is a parameter to control history
  - can also use queue length or buffer utilization
- When utilization is above a threshold
  - for each new packet to be sent over congested link
    - send “choke” packet back to sender
    - tag forwarded data packet to prevent more coke packets
  - when sender receives choke packet
    - must reduce rate to “choked” destination
- Hop-by-hop coke
  - on path back to sender, each router reduces traffic
  - consumes buffer space along path to sender
  - provides faster relief to congested router/link

Fair Queuing

- Local (per router) congestion control
  - each output link has n queues, one for each sender
    - need to limit max queue size or buffers will be exhausted
  - use round-robin to select next packet to queue
    - can use per-packet or per-byte

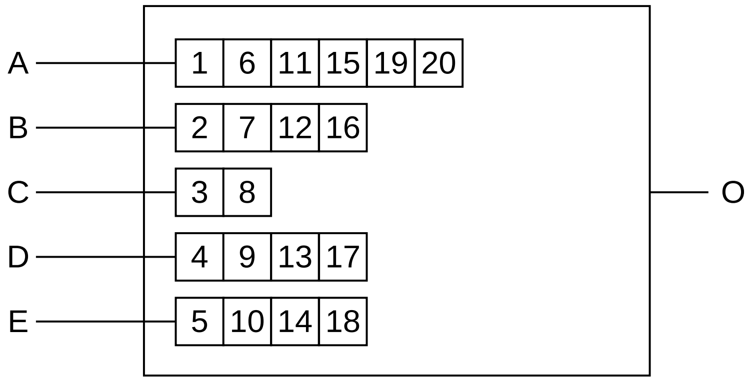

(a)

Packet	Finishing time
C	8
B	16
D	17
E	18
A	20

(b)

From: *Computer Networks*, 3rd Ed. by Andrew S. Tanenbaum, (c)1996 Prentice Hall.

- Weighted Fair Queuing
  - can give different links different priorities
  - give higher priority length multiple slots per round

Load Shedding

- When all else fails, routers drop (discard) packets
- Policy question: what packets to drop?
  - oldest ones: they are likely to be useless now
  - newest ones: helps to close open window in file transfer
  - less important ones
    - requires cooperation of application
    - in MPEG I frames are more important than B frames
  - drop all related packets
    - fragmentation: loss of one packet renders others useless
    - requires information from higher levels
- Preemptive shedding
  - when traffic starts to get high, dropping packets can prevent additional congestion

RSVP - Multicast Bandwidth Reservation

- Receivers send request to reserve BW up spanning tree
- Routers propagate request if request up tree
  - only sent if greater than prev. request for this group
- Dest. can request BW for multiple alternative sources
  - routers only allocate bandwidth for maximum channel request

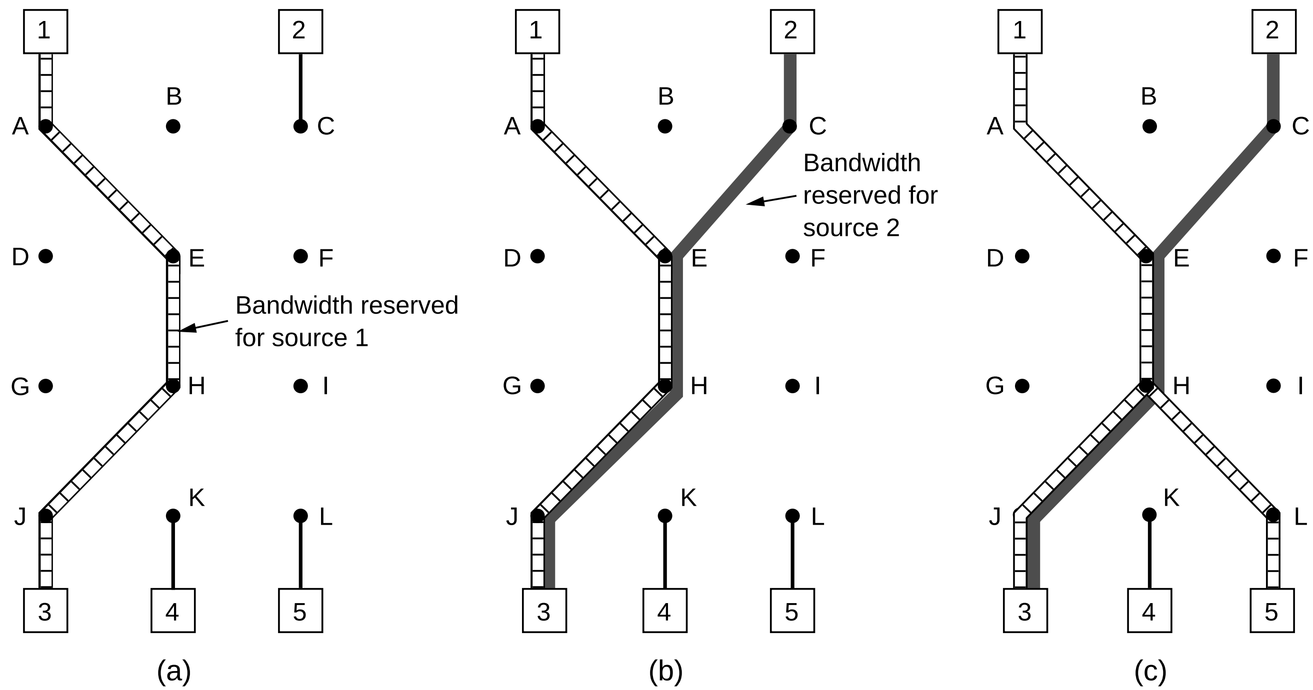

From: *Computer Networks*, 3rd Ed. by Andrew S. Tanenbaum, (c)1996 Prentice Hall.