

CONGESTION CONTROL

Phenomenon: when too much traffic enters into system, performance degrades

—→ excessive traffic can cause congestion



Problem: regulate traffic influx such that congestion does not occur


—→ not too fast, not too slow

—→ congestion control

—→ first question: what is congestion?

Viewpoint: 3 components

→ (1) traffic coming in, (2) in transit, (3) going out



At time instance t :

- traffic influx: $\lambda(t)$ “offered load” (bps)
- traffic outflux: $\gamma(t)$ “throughput” (bps)
- traffic in-flight: $Q(t)$ “load” (volume, i.e., no. of packets)

Examples:

Highway system:

- traffic influx: no. of cars entering highway per second
- traffic outflux: no. of cars exiting highway per second
- traffic in-flight: no. of cars traveling on highway

→ at time instance t



California Dept. of Transportation (Caltrans)

Water faucet and sink:

- traffic influx: water influx per second
- traffic outflux: water outflux per second
- traffic in-flight: water level in sink

→ not good if sink overflows



faucet.com

Many examples: heating/cooling system with thermostat, car cruise control, ...

What is the meaning of congestion?

→ when sending too fast, throughput starts to go down

In the water faucet/sink example: is there congestion?

What about highway system?

What we can control:

→ traffic influx rate $\lambda(t)$

→ no power over anything else

Congestion control: how to regulate influx rate $\lambda(t)$ —not too fast, not too slow—so that throughput $\gamma(t)$ is maximized

→ many applications

→ TCP congestion control

→ multimedia video/audio streaming

Pseudo real-time multimedia streaming:

Examples: streaming client/server apps

→ real-time vs. pseudo-real-time

“Pseudo” because of prefetching trick

→ application is given headstart before playback

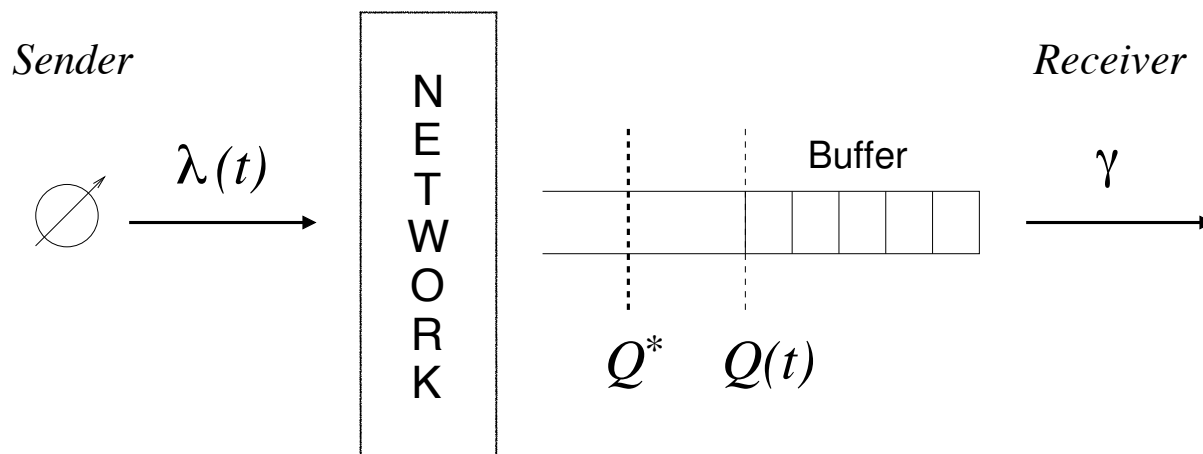
→ fill & prevent client buffer from becoming empty

Main steps:

- prefetch X seconds worth of audio/video data
 - initial playback delay
- keep fetching audio/video data such that X seconds worth of future data resides in receiver's buffer
 - protects against spurious congestion
 - don't keep more than X
 - potential waste of resources: bandwidth, memory, CPU

If streaming is done well, user experiences continuous playback without quality disruptions

Pseudo real-time application architecture:



- $Q(t)$: current buffer level
- Q^* : desired buffer level
- γ : throughput—fixed playback rate
→ e.g., 24 frames-per-second (fps) for movies

Goal: keep $Q(t) \approx Q^*$ by adjusting $\lambda(t)$

- don't buffer too much: resource wastage
- don't buffer too little: cannot hide congestion

How does load $Q(t)$ vary?

→ obeys simple rule

Compare two time instances t and $t + 1$.

At time $t + 1$:

$$Q(t + 1) = Q(t) + \lambda(t) - \gamma(t)$$

- $Q(t)$: what was there to begin with
- $\lambda(t)$: what newly arrived
- $\gamma(t)$: what newly exited
- $\lambda(t) - \gamma(t)$: net influx (positive or negative)
- note: $Q(t)$ cannot be negative by its meaning
 - no. of packets
 - $Q(t + 1) = \max\{0, Q(t) + \lambda(t) - \gamma(t)\}$
- missing item?

Other applications.

Ex. 1: Router congestion control

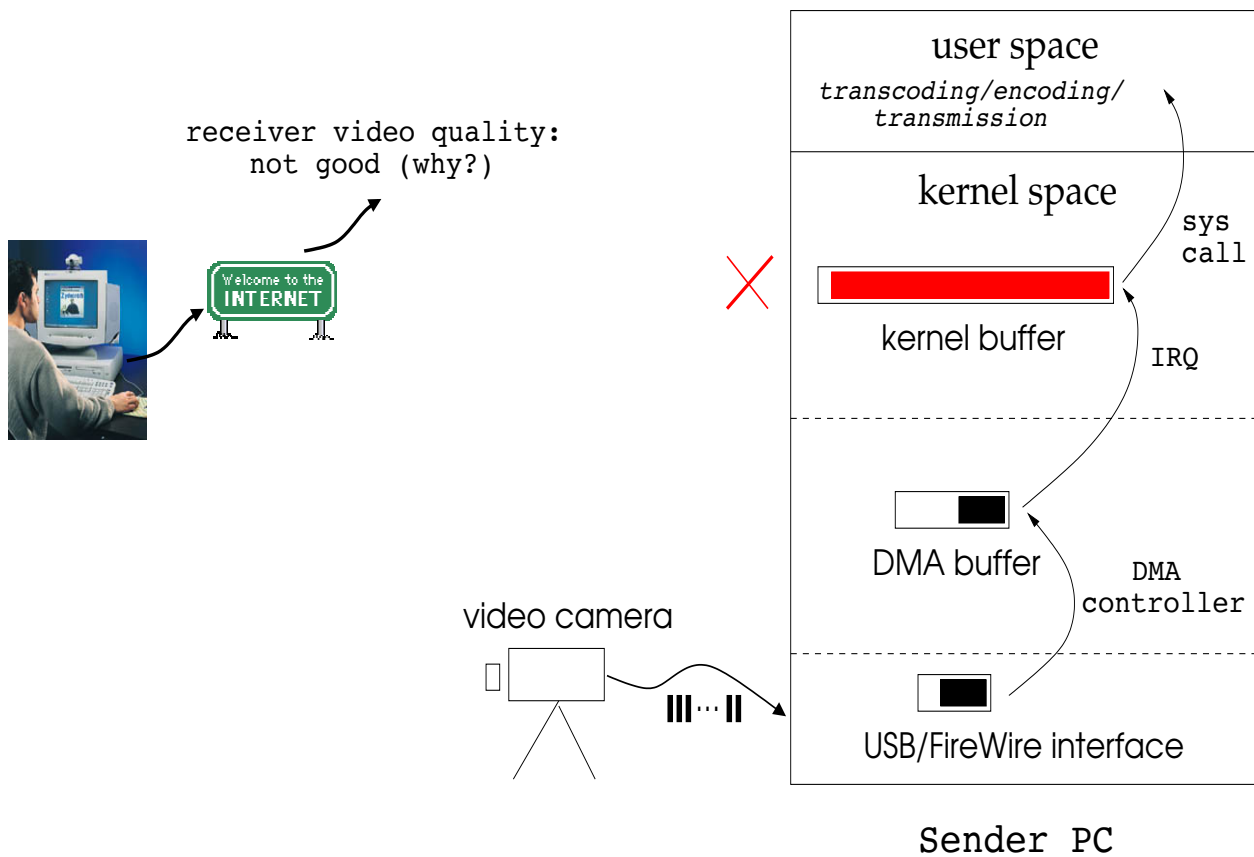
→ active queue management (AQM)

- receiver is a router/switch
- Q^* is desired buffer occupancy/delay at router
 - too much buffering: bufferbloat
- router throttles sender(s) to maintain Q^*
 - router sends control packets to senders
 - instruction: slow down, go faster, stay put

Ex. 2: Desktop videoconferencing

→ video quality may not be good: why?

→ common misconception: sole culprit is network



What is the goal:

- achieve $Q(t) = Q^*$
- or close to it: $|Q(t) - Q^*| < \varepsilon$

Basic idea:

- if $Q(t) = Q^*$ do nothing
- if $Q(t) < Q^*$ increase $\lambda(t)$
 - too little in the buffer
- if $Q(t) > Q^*$ decrease $\lambda(t)$
 - too much in the buffer

Rule of thumb: called control law

Since state of receiver buffer must be conveyed to sender who adjusts $\lambda(t)$:

- called feedback control
- also closed-loop control

Network protocol implementation:

→ design choices

- control action undertaken at sender

→ smart sender/dump receiver

→ preferred mode of Internet protocols

→ when might the opposite be better?

- receiver informs sender of Q^* and $Q(t)$

→ feedback could just be gap $Q^* - Q(t)$

→ or simply up/down binary indication

Key question in feedback congestion control:

→ how to increase/decrease $\lambda(t)$

Desired state of the system:

$$Q(t) = Q^* \text{ and } \lambda(t) = \gamma$$

→ why is $\lambda(t) = \gamma$ needed?

→ system is in equilibrium or steady-state



Starting state:

→ empty buffer and nothing is being sent

→ think of iTunes, Netflix, Spotify, etc.

$$\text{i.e., } Q(t) = 0 \text{ and } \lambda(t) = 0$$

Time evolution (or dynamics): track $Q(t)$ and $\lambda(t)$



Congestion control methods: A, B, C and D

Method A:

- if $Q(t) = Q^*$ then $\lambda(t + 1) \leftarrow \lambda(t)$
- if $Q(t) < Q^*$ then $\lambda(t + 1) \leftarrow \lambda(t) + a$
- if $Q(t) > Q^*$ then $\lambda(t + 1) \leftarrow \lambda(t) - a$

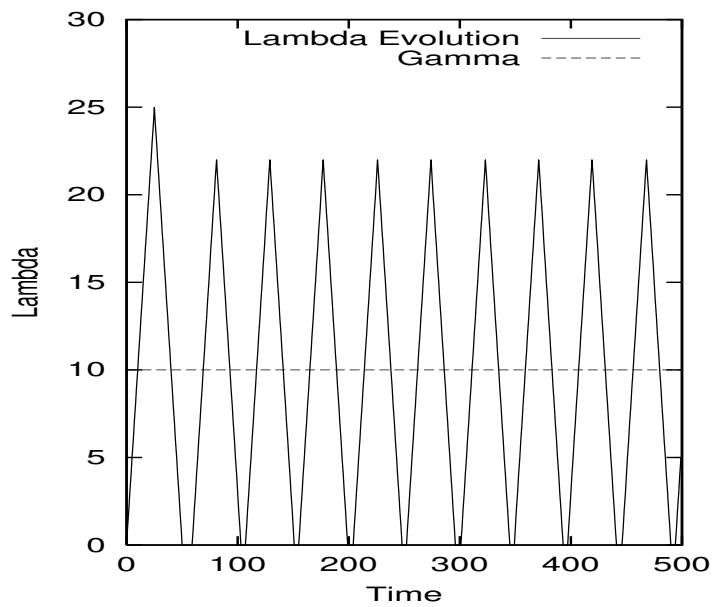

where $a > 0$ is a fixed parameter

→ called linear increase/linear decrease



Question: how well does it work?

Example:


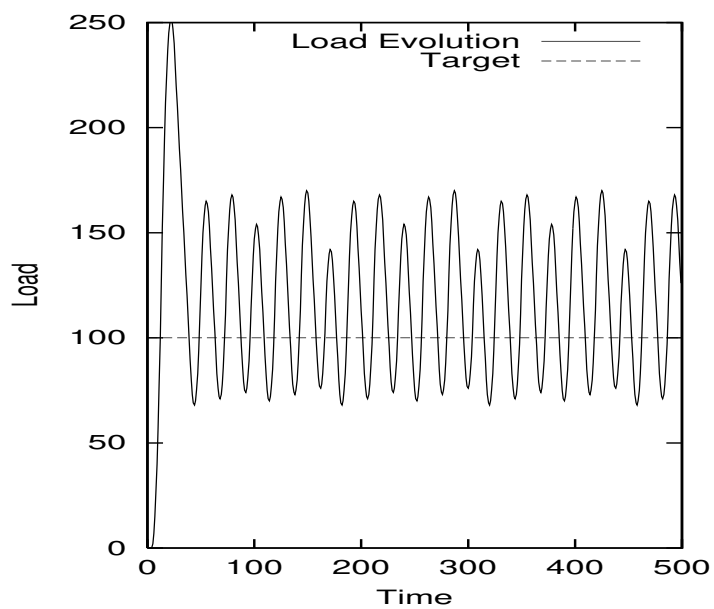
- $Q(0) = 0$
- $\lambda(0) = 0$
- $Q^* = 100$
- $\gamma = 10$
- $a = 1$



With $a = 0.5$:





With $a = 3$:



Remarks:

- Method A isn't that great no matter what a value is used
 - keeps oscillating
- Actually: would lead to unbounded oscillation if not for physical restriction $\lambda(t) \geq 0$ and $Q(t) \geq 0$
 - i.e., bottoms out
 - easily seen: start from non-zero buffer
 - e.g., $Q(0) = 110$

With $a = 1$, $Q(0) = 110$, $\lambda(0) = 11$:



Method B:

- if $Q(t) = Q^*$ then $\lambda(t + 1) \leftarrow \lambda(t)$
- if $Q(t) < Q^*$ then $\lambda(t + 1) \leftarrow \lambda(t) + a$
- if $Q(t) > Q^*$ then $\lambda(t + 1) \leftarrow \delta \lambda(t)$

where $a > 0$ and $0 < \delta < 1$ are fixed parameters



Note: only decrease part differs from **Method A**.

- linear increase with slope a
- exponential decrease with backoff factor δ
- e.g., binary backoff in case $\delta = 1/2$



Similar to Ethernet and WLAN backoff

- question: does it work?



With $a = 1$, $\delta = 1/2$:



With $a = 3$, $\delta = 1/2$:



With $a = 1$, $\delta = 3/4$:



Note:

- Method B oscillates around target
 - does not converge
- Superior to Method A: unbounded oscillation
 - doesn't hit "rock bottom"
 - due to asymmetry in increase vs. decrease policy
 - we observe "sawtooth" pattern
- Method B is used by TCP
 - linear increase/exponential decrease
 - additive increase/multiplicative decrease (AIMD)
 - will discuss separately under TCP congestion control

Question: can we do better?

→ what information have we not made use of?

Method C:

$$\lambda(t + 1) \leftarrow \lambda(t) + \varepsilon(Q^* - Q(t))$$

where $\varepsilon > 0$ is a fixed parameter

Tries to adjust magnitude of change as a function of the gap $Q^* - Q(t)$

→ if gap is big, change by a lot

→ if gap is small, change by a little

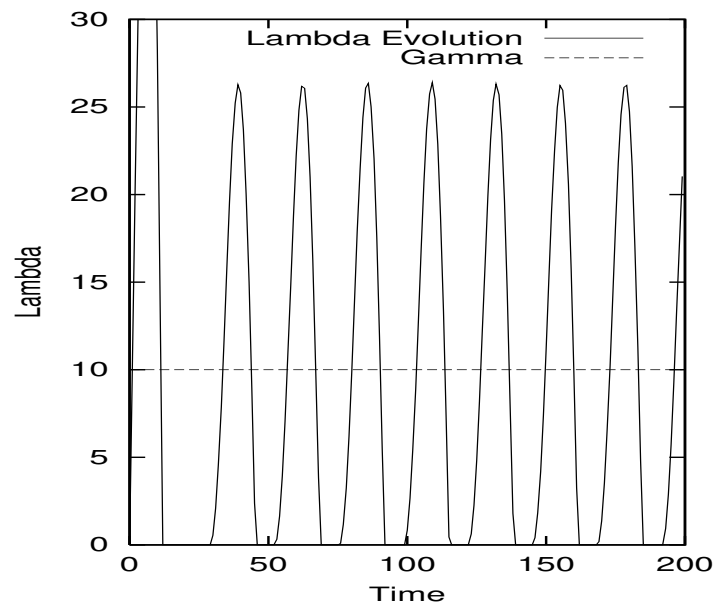

Thus:

- if $Q^* - Q(t) > 0$, increase $\lambda(t)$ proportional to gap
- if $Q^* - Q(t) < 0$, decrease $\lambda(t)$ proportional to gap



Trying to be more clever...

→ bottom line: is it any good?

With $\varepsilon = 0.1$:



With $\epsilon = 0.5$:



Answer: no

→ control law looks good on the surface

→ but looks can be deceiving

Time to try something different.

Method D:

$$\lambda(t + 1) \leftarrow \lambda(t) + \varepsilon(Q^* - Q(t)) + \beta(\gamma - \lambda(t))$$



where $\varepsilon > 0$ and $\beta > 0$ are fixed parameters

→ additional term $\beta(\gamma - \lambda(t))$

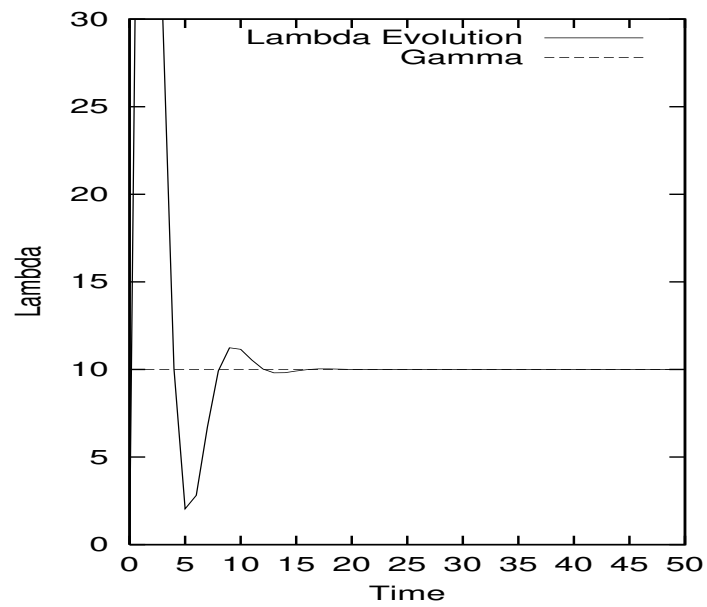

→ what's going on?

→ does it work?


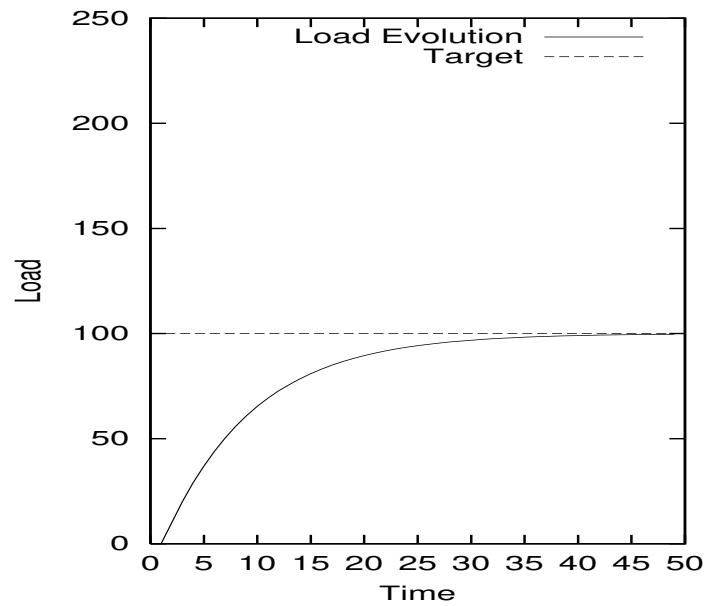
With $\varepsilon = 0.2$ and $\beta = 0.5$:



With $\varepsilon = 0.5$ and $\beta = 1.1$:



With $\varepsilon = 0.1$ and $\beta = 1.0$:



Remarks:

- Method D has desired behavior
- Superior to Methods A, B, and C,
- No unbounded oscillation
- Convergence to desired state
 - target operating point (Q^*, γ)
 - asymptotically stable

Related to PID control of physical systems.

→ e.g., drones