

What is the goal?

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

How to: 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
 - a 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:

- some design options available
- control action undertaken at sender
 - smart sender/dumb receiver
 - preferred mode of many Internet protocols
 - when might the opposite be better?
- receiver informs sender of Q^* and $Q(t)$
 - feedback packet (control signaling/messaging)
 - feedback could just be gap $Q^* - Q(t)$
 - or simply up/down binary indication

Key question in feedback congestion control:

- **how much** to increase/decrease $\lambda(t)$
- already know in which direction

Desired state of the system:

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



- why is “ $\lambda(t) = \gamma$ ” needed?

Starting state:

- empty buffer and nothing is being sent
- think of iTunes, Rhapsody, 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:

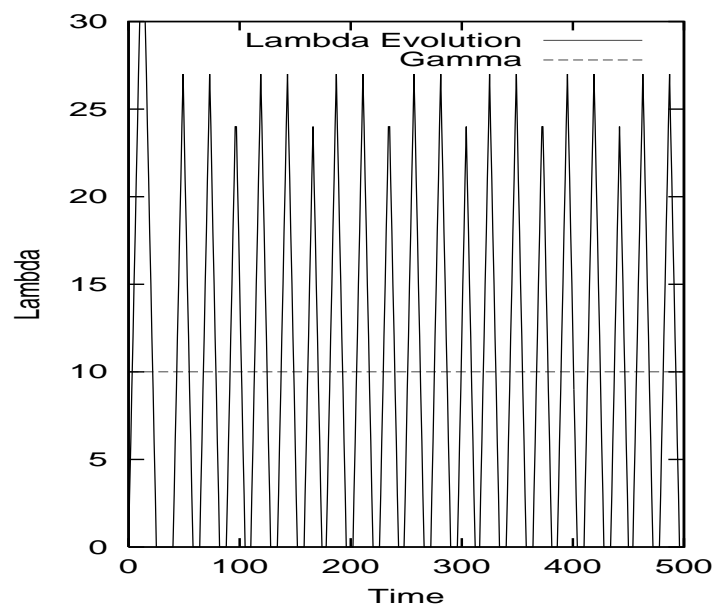

- $Q(0) = 0$
- $\lambda(0) = 0$
- $Q^* = 100$
- $\gamma = 10$
- $a = 1$



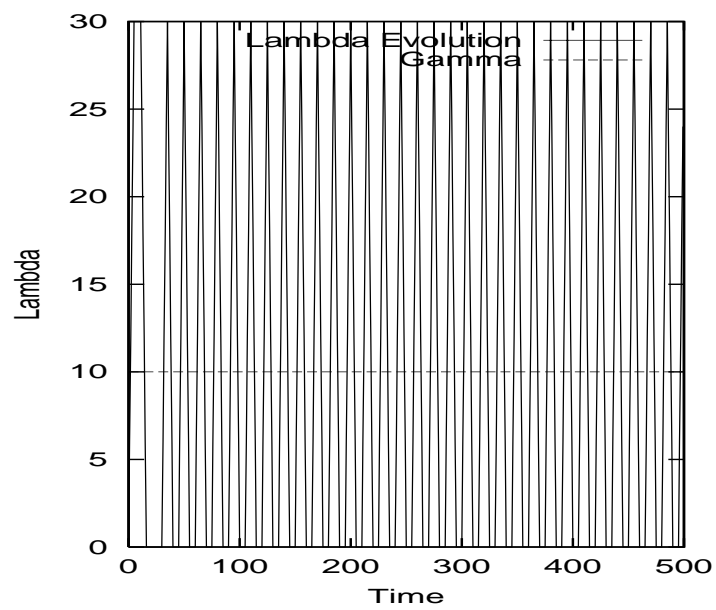

With $a = 0.5$:



With $a = 3$:





With $a = 6$:



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 $\lambda(t + 1) \leftarrow \delta \cdot \lambda(t)$

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



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



Note:

- Method B isn't that great either
- One advantage over Method A: doesn't lead to unbounded oscillation
 - note: 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)

Question: can we do better?

→ what "freebie" 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 $Q^* - Q(t)$

→ incorporate distance from target Q^*

→ before: just the sign (above/below)



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 $\varepsilon = 0.5$:



Answer: no

→ control law looks good on the surface

→ but looks can be deceiving

Time to try something strange

→ any (crazy) ideas?