Congestion Control

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

$\longrightarrow$ excessive traffic can cause congestion

Problem: regulate traffic influx such that congestion does not occur

$\longrightarrow$ congestion control

Need to understand:

- What is congestion?
- How do we prevent or manage it?

Traffic influx/outflux picture:

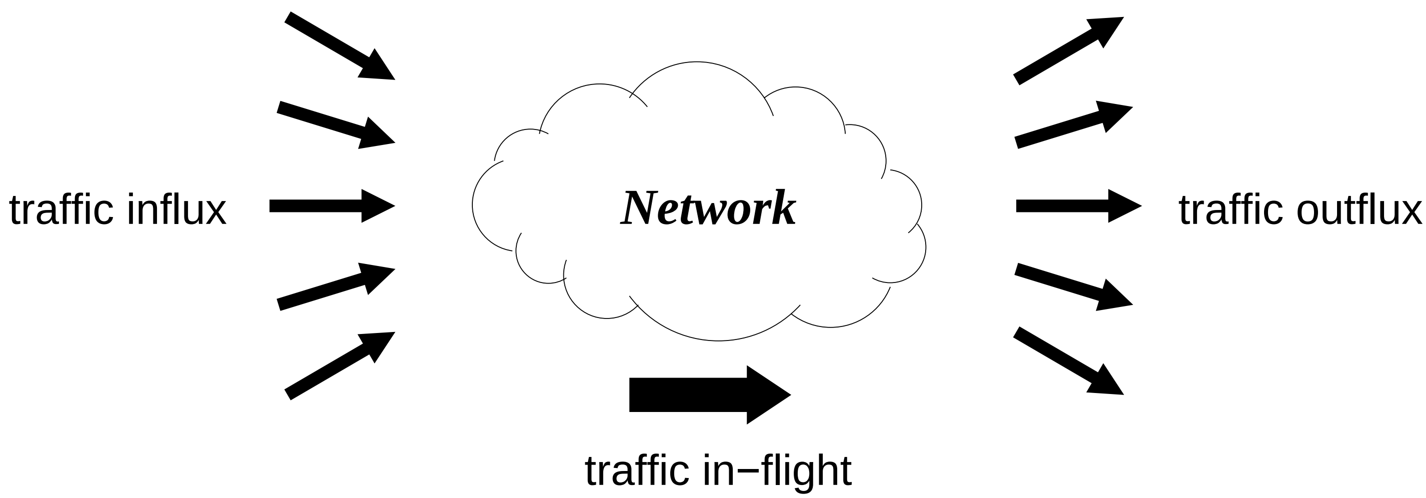

- traffic influx: $\lambda(t)$ "offered load"

  $\rightarrow$ rate: bps (or pps) at time t

- traffic outflux: $\gamma(t)$ "throughput"

  $\rightarrow$ rate: bps (or pps) at time t

- traffic in-flight: $Q(t)$

  $\rightarrow$ volume: total packets in transit at time t

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

$\longrightarrow$ 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

$\longrightarrow$ "congestion?"

faucet.com

Thermostat . . .

802.11b WLAN:

- Throughput

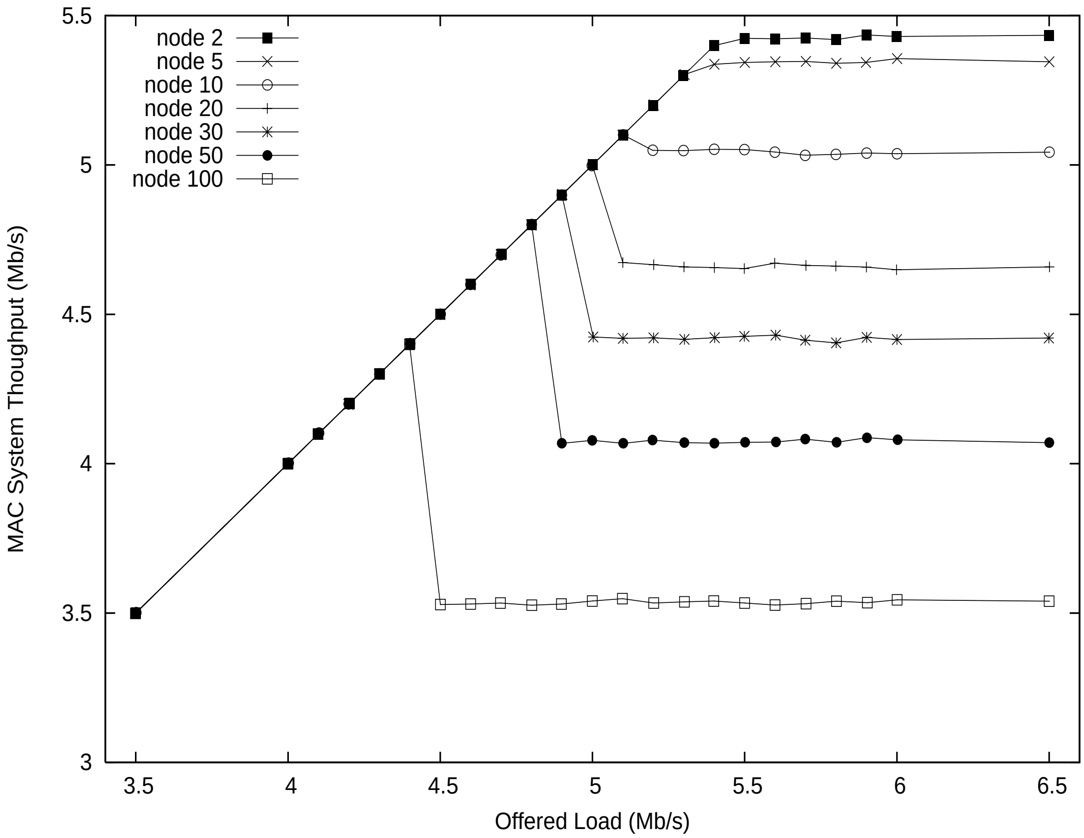

$\longrightarrow$ unimodal or bell-shaped

$\longrightarrow$ recall: less pronounced in real systems

802.11b WLAN:

- Collision

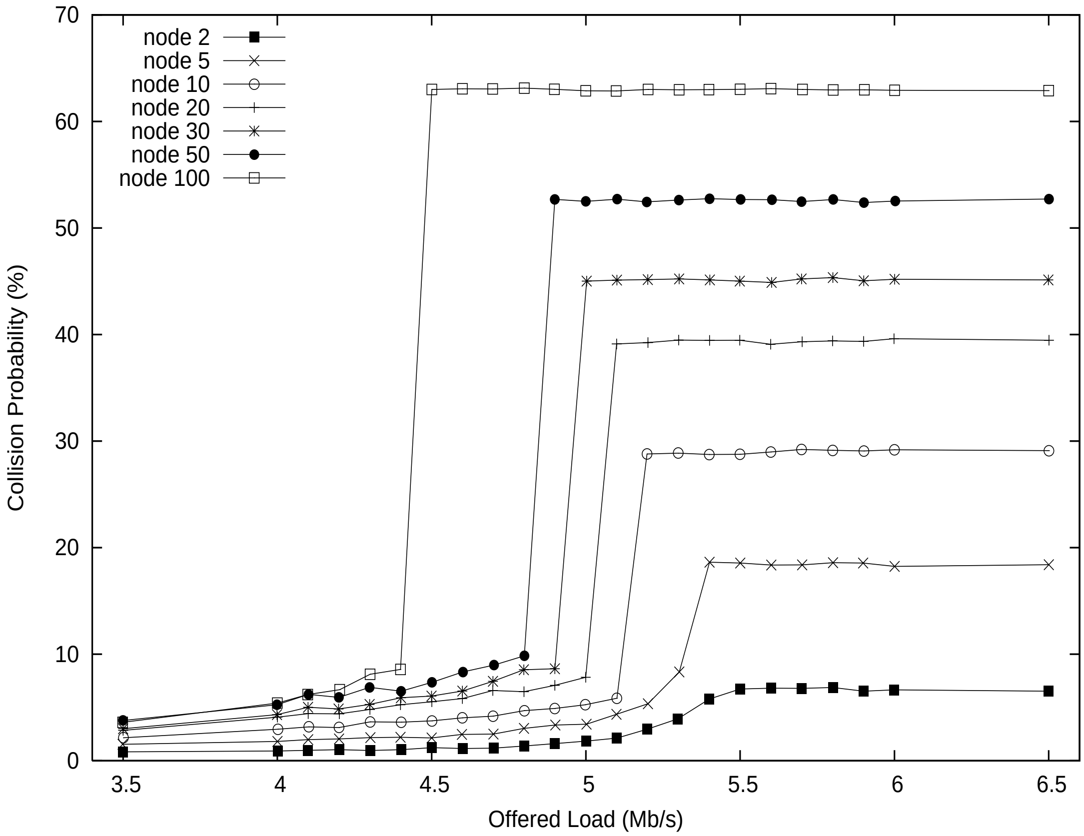

$\longrightarrow$ underlying cause of unimodal throughput

What we can regulate or control:

$\longrightarrow$ traffic influx rate $\lambda(t)$

Ex.:

- Faucet knob in water sink
- Temperature needle in thermostat
- Cars entering onto highway
- Traffic sent by UDP or TCP

What we cannot control: the rest

$\longrightarrow$ except in the long run: bandwidth planning

$\longrightarrow$ does scheduling help?

$\longrightarrow$ Kleinrock's conservation law: "zero-sum pie"

How does in-flight traffic or load $Q(t)$ vary?

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 (delivered to applications)
- $\lambda(t) - \gamma(t)$: net influx
- $Q(t)$ cannot be negative

  $\longrightarrow Q(t+1) = \max\{0, Q(t) + \lambda(t) - \gamma(t)\}$
- missing item?

Goal: Want to keep system in "good/desirable" state

Ex.: If $\lambda(t) > \gamma(t)$ for all time then

$$Q(t) \rightarrow \infty \quad \text{as} \quad t \rightarrow \infty$$

$\longrightarrow$ water level in sink grows and grows

$\longrightarrow$ water sink has finite "buffer" capacity, overflows

$\longrightarrow$ want to keep water level stable; how?

Control actions:

- If water level is too high, close faucet
- If water level is too low, open faucet

$\longrightarrow$ feedback control

$\longrightarrow$ "state of system": water level

Pseudo Real-Time Multimedia Streaming

$\longrightarrow$ e.g., RealPlayer, Rhapsody, Internet radio

$\longrightarrow$ "pseudo" because of prefetching trick

$\longrightarrow$ application is given headstart: few seconds

$\longrightarrow$ why?

Goal: fill buffer & prevent from becoming empty

Method:

- prefetch X seconds worth of data (e.g., audio/video)
- initial delayed playback: penalty of pseudo real-time
- keep fetching audio/video data such that X seconds worth of future data resides in receiver's buffer

  $\rightarrow$ allows hiding of spurious congestion

  $\rightarrow$ user: continuous playback experience

  $\rightarrow$ can it work if bandwidth $<$ app data rate?

Pseudo real-time traffic control architecture:

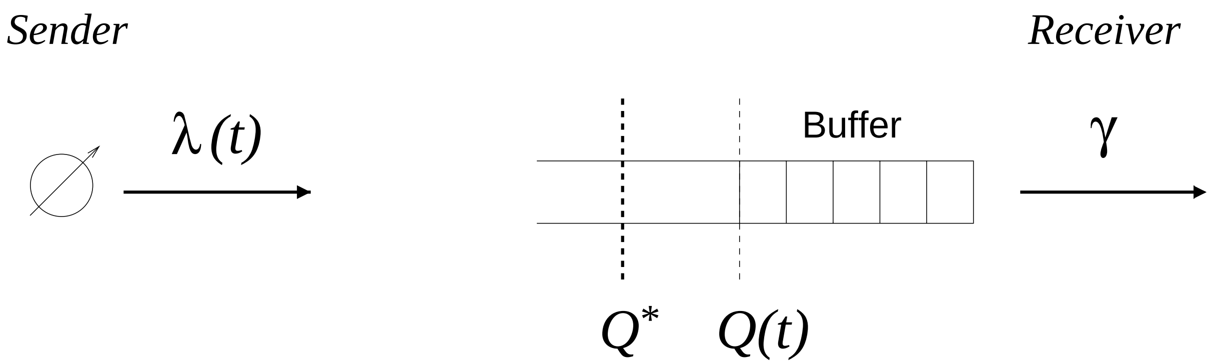

- $Q(t)$: current buffer level
- Q^*: desired buffer level
- γ: throughput, i.e., playback rate

  $\rightarrow$ e.g., for video 24 frames-per-second (fps)

Goal: vary $\lambda(t)$ such that $Q(t) \approx Q^*$

$\longrightarrow$ don't buffer too much (memory cost)

$\longrightarrow$ don't buffer too little (bumpy road)

Basic idea:

- if $Q(t) = Q^*$ do nothing
- if $Q(t) < Q^*$ increase $\lambda(t)$
- if $Q(t) > Q^*$ decrease $\lambda(t)$

$\longrightarrow$ "control law"

$\longrightarrow$ thermostat control (same as water faucet)

Protocol implementation:

- control action undertaken at sender

  $\rightarrow$ smart sender/dump receiver

  $\rightarrow$ when might the opposite be better?

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

  $\rightarrow$ feedback packet ("control signaling")

  $\rightarrow$ or just $Q^* - Q(t)$

  $\rightarrow$ or just up/down (binary)

Other applications:

Router congestion control

$\longrightarrow$ active queue management (AQM)

- receiver is a router
- Q^* is desired buffer occupancy/delay at router
- router throttles sender(s) to maintain Q^*

$\longrightarrow$ similar to old source quench message (ICMP)

$\longrightarrow$ considered too much messaging overhead

Slightly modified Internet standard:

$\longrightarrow$ ECN (explicit congestion notification)

- two bits in IPv4 TOS field

  $\rightarrow$ ECT: ECN capable transport (bit 6)

  $\rightarrow$ CE: congestion experienced (bit 7)

- congested router marks ECT
- supported in most routers, default not turned on
- requires TCP sender/receiver changes

Also proposed to throttle denial-of-service attack traffic

$\longrightarrow$ push-back

$\longrightarrow$ good guy vs. bad guy problem

Key question in feedback congestion control: how much to increase or decrease $\lambda(t)$

$\longrightarrow$ "control problem"

$\longrightarrow$ different specific manifestation

$\longrightarrow$ TCP has its own specific rule

Desired state of the system:

$\longrightarrow$ i.e., target operating point

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

Start from:

$\longrightarrow$ empty buffer and no sending rate at start

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

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

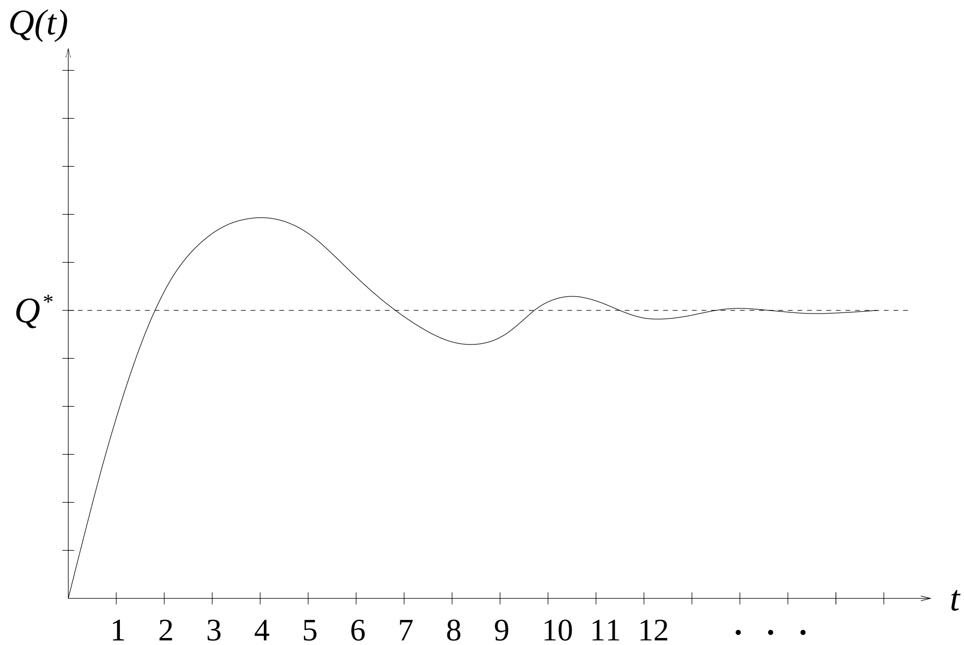

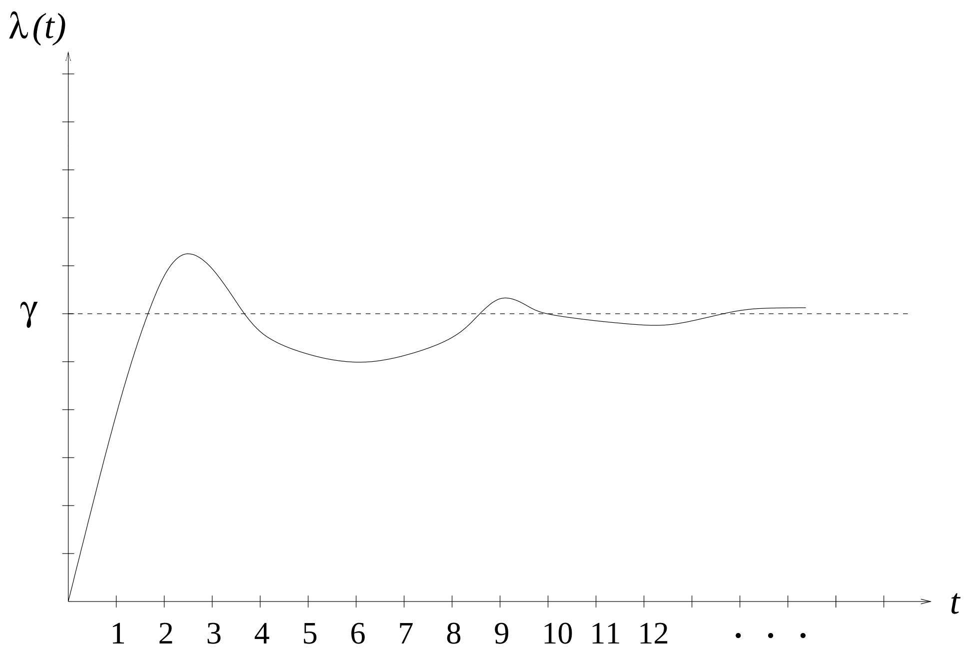

Congestion control methods: A, B, C and D

`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

$\longrightarrow$ linear increase and linear decrease

Question: does it work?

Example:

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

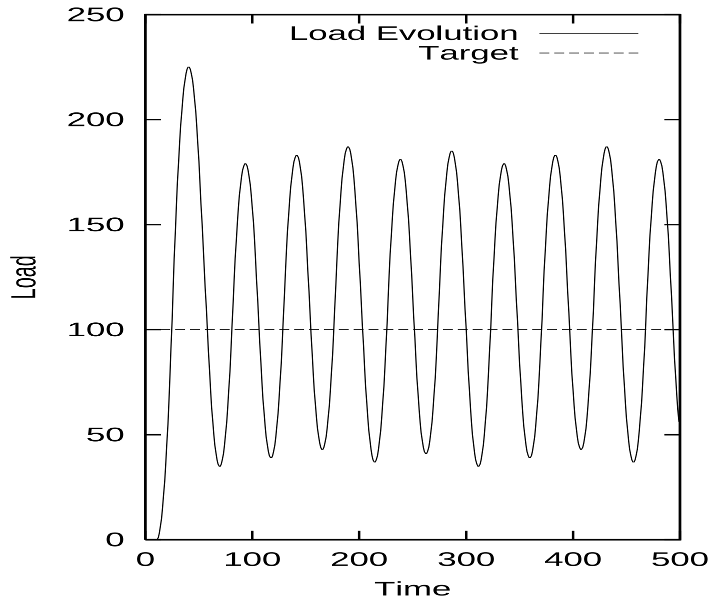
Load Evolution
Target
Load
Time
250
200
150
100
50
0
0
100
200
300
400
500

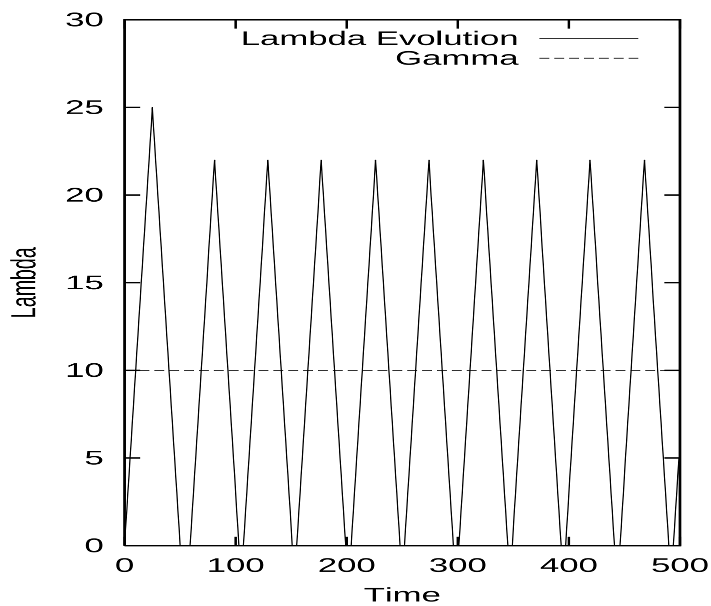
Lambda Evolution
Gamma
Lambda
Time
30
25
20
15
10
5
0
0
100
200
300
400
500

With $a = 0.5$:

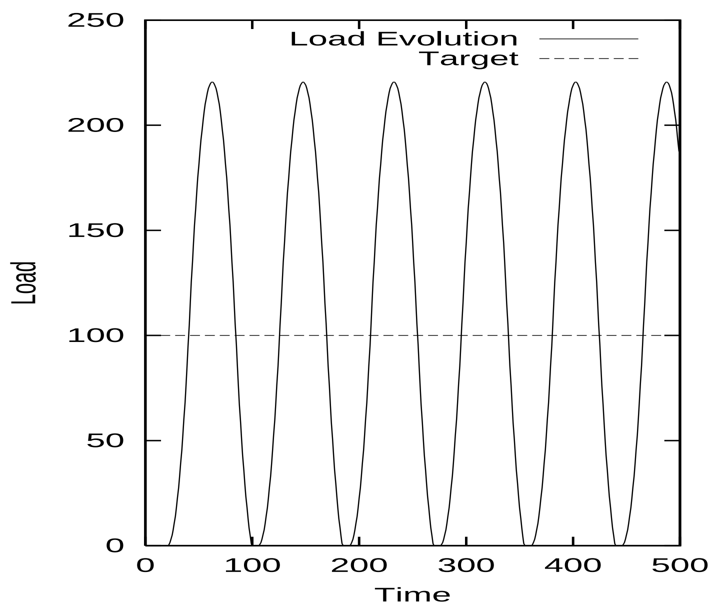
Load Evolution
Target
Load
Time
250
200
150
100
50
0
0
100
200
300
400
500

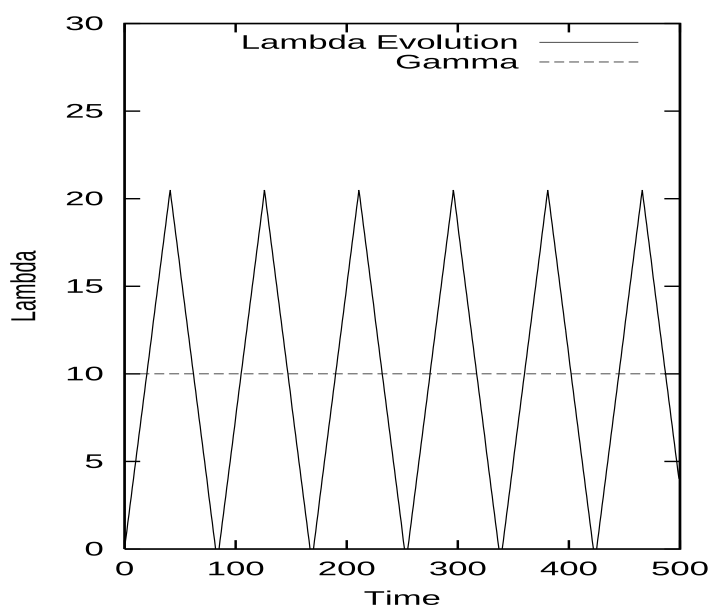
Lambda Evolution
Gamma
Lambda
Time
30
25
20
15
10
5
0
0
100
200
300
400
500

With $a = 3$:

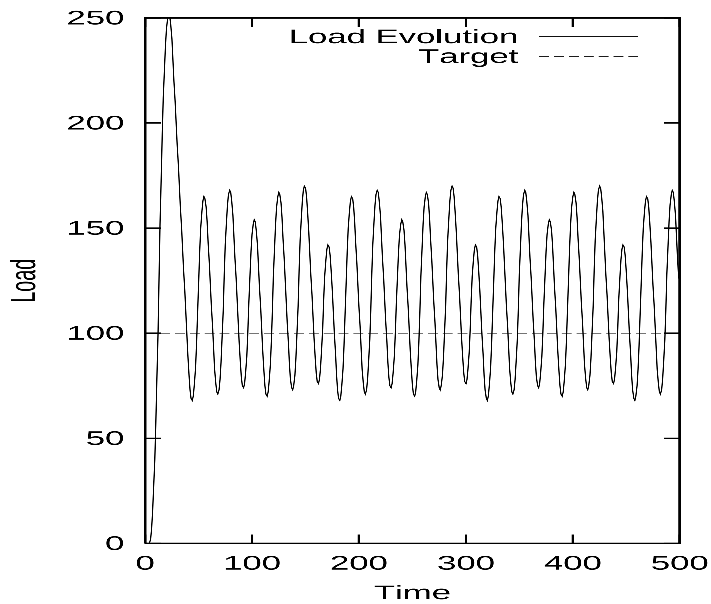

Load Evolution
Target
250
200
150
100
50
0
Load
0
100
200
300
400
500
Time

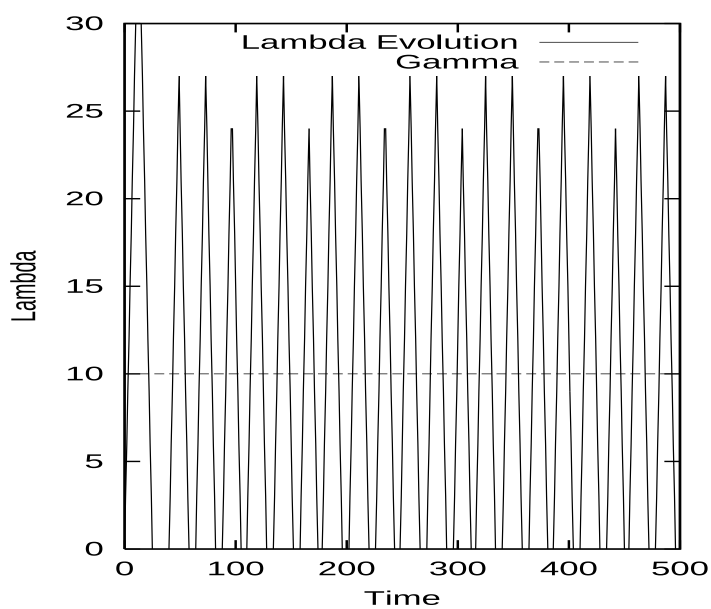

Lambda Evolution
Gamma
30
25
20
15
10
5
0
Lambda
0
100
200
300
400
500
Time

With $a = 6$:

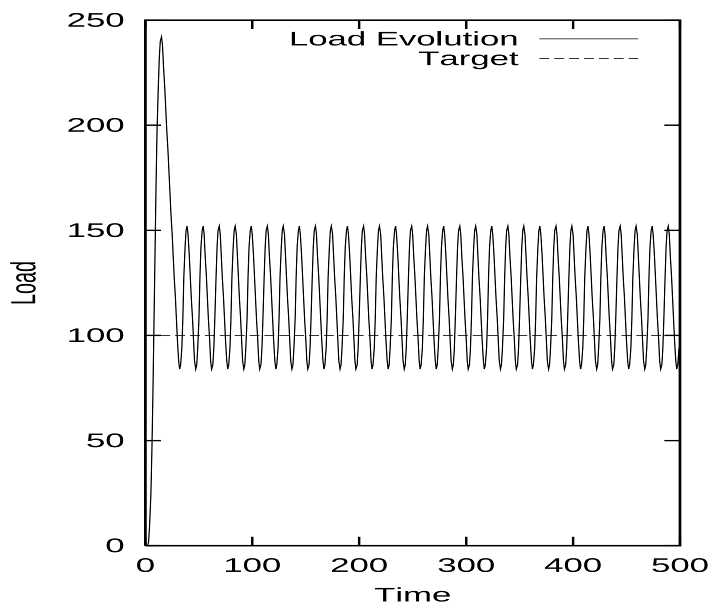
Load Evolution
Target
Load
Time
0
50
100
150
200
250
0
100
200
300
400
500

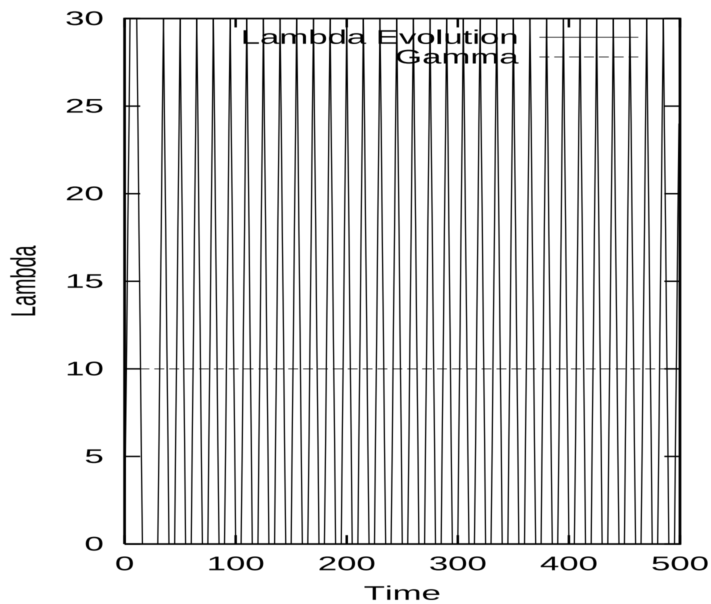
Lambda Evolution
Gamma
Lambda
Time
0
5
10
15
20
25
30
0
100
200
300
400
500

Remarks:

- Method A isn't that great no matter what a value is used

  $\rightarrow$ keeps oscillating

- Actually: would lead to unbounded oscillation if not for physical restriction $\lambda(t) \geq 0$ and $Q(t) \geq 0$

  $\longrightarrow$ easily seen: start from non-zero buffer

  $\longrightarrow$ e.g., $Q(0) = 110$

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

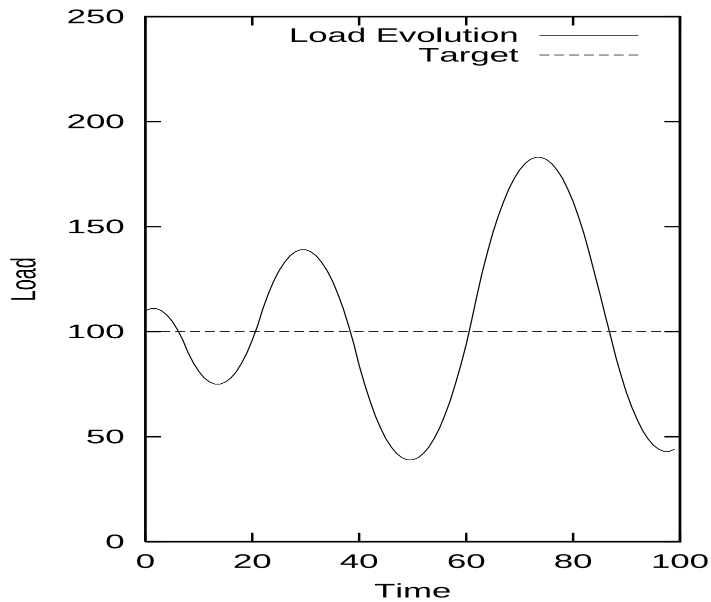

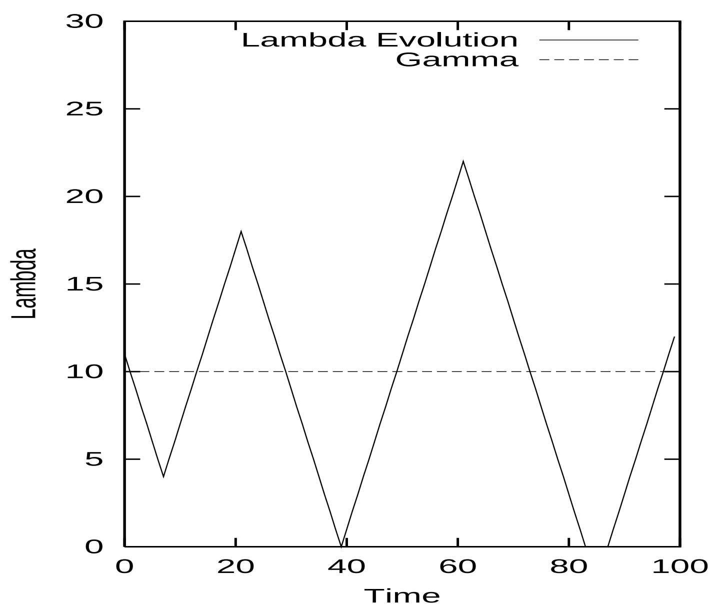

`Method B`:

- 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`.

$\longrightarrow$ linear increase with slope a

$\longrightarrow$ exponential decrease with backoff factor δ

$\longrightarrow$ e.g., binary backoff in case $\delta = 1/2$

Similar to Ethernet and WLAN backoff

$\longrightarrow$ question: does it work?

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

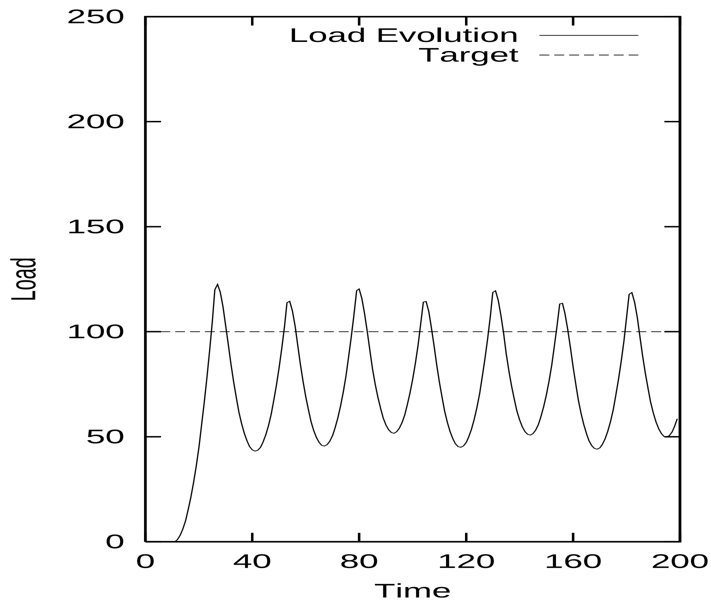

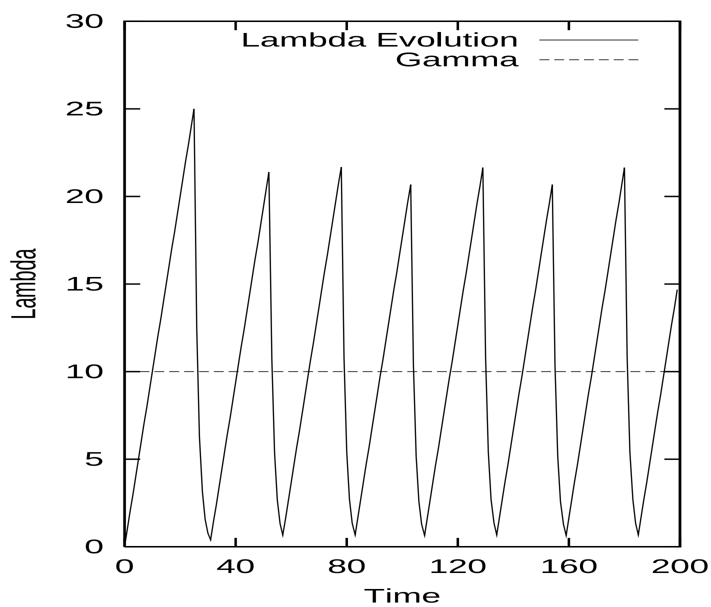

With $a = 3$, $\delta = 1/2$:

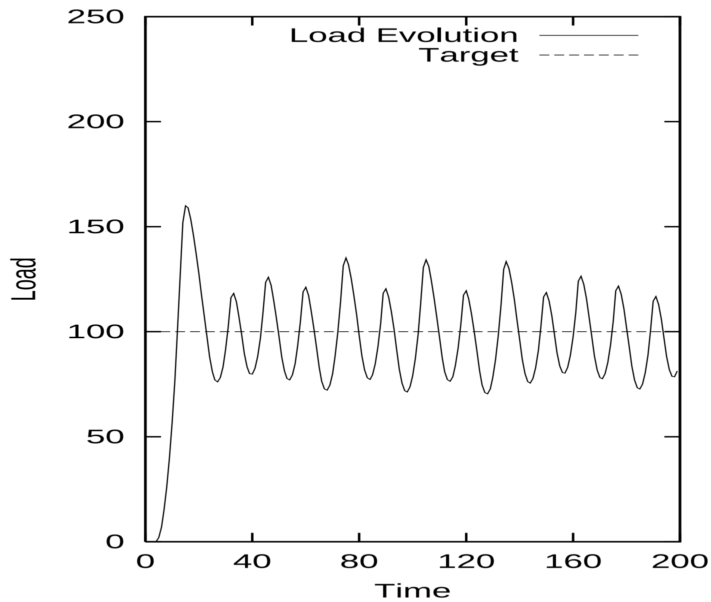

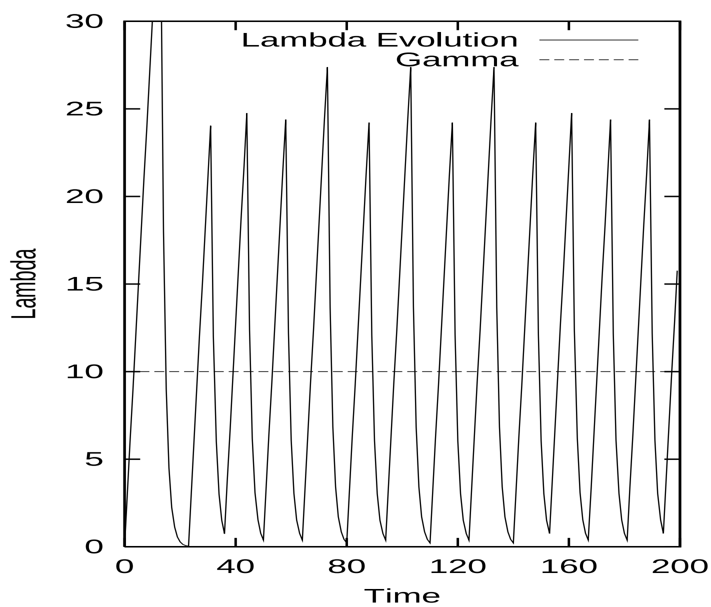

With $a = 1$, $\delta = 1/4$:

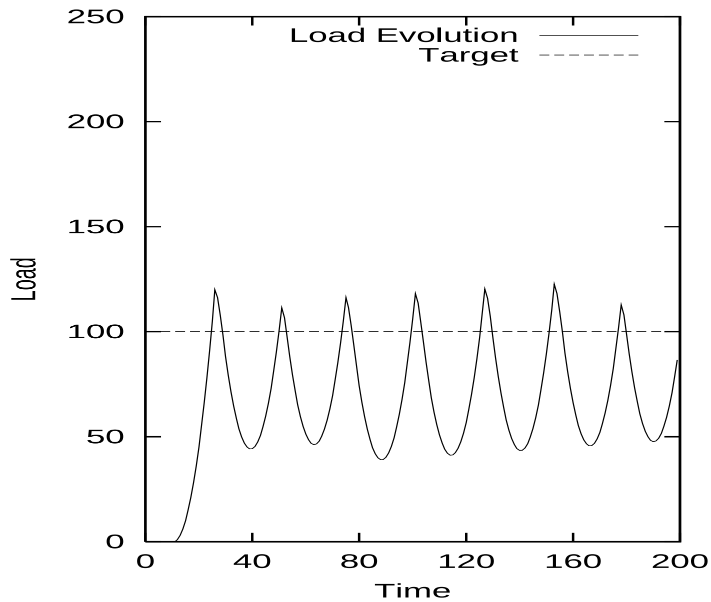

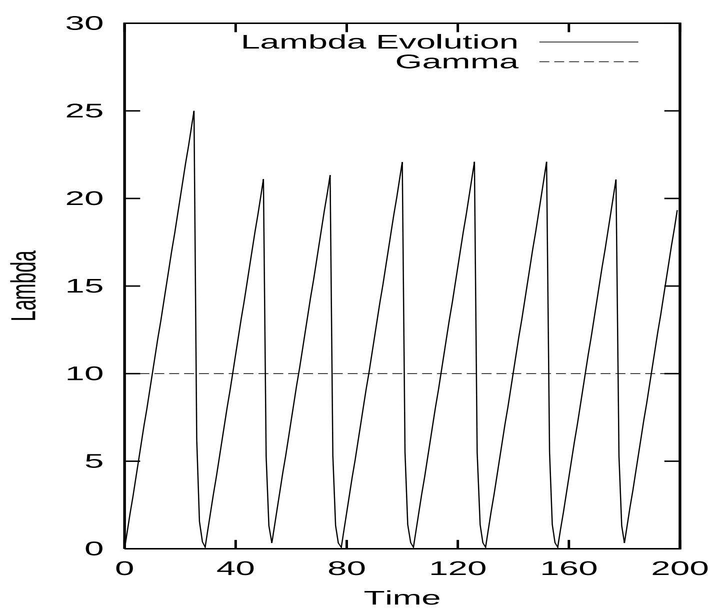

With $a = 1$, $\delta = 3/4$:

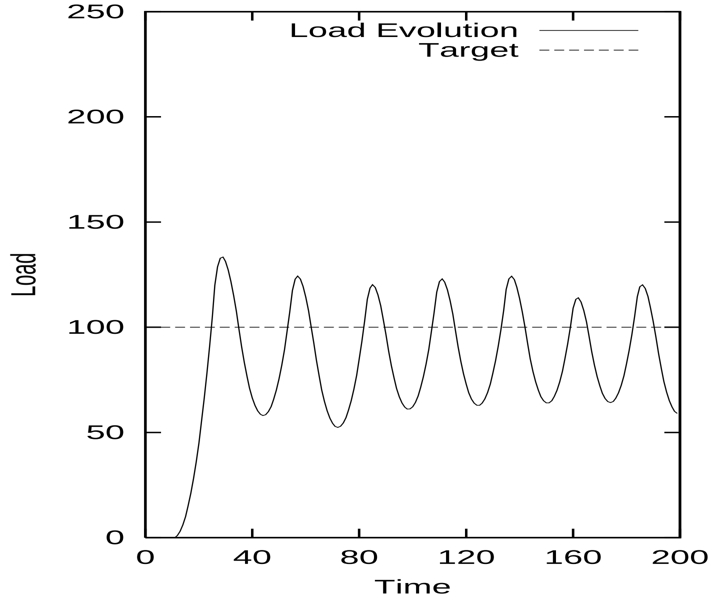

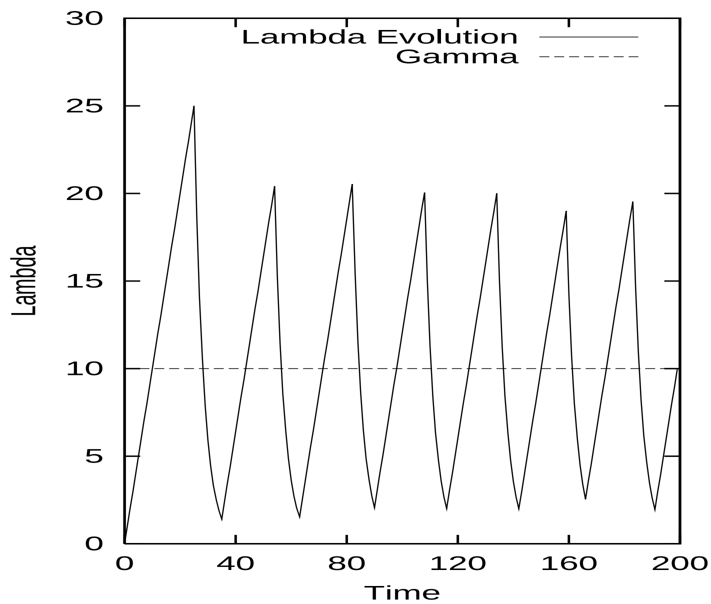

Note:

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

  $\rightarrow$ note: doesn't hit "rock bottom"

  $\rightarrow$ due to asymmetry in increase vs. decrease policy

  $\rightarrow$ typical "sawtooth" pattern
- Method B is used by TCP

  $\rightarrow$ linear increase/exponential decrease

  $\rightarrow$ additive increase/multiplicative decrease (AIMD)

Question: can we do better?

$\longrightarrow$ what "freebie" have we not utilized yet?

`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 $Q^* - Q(t)$

$\longrightarrow$ incorporate distance from target Q^*

$\longrightarrow$ 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. . .

$\longrightarrow$ bottom line: is it any good?

With $\varepsilon = 0.1$:

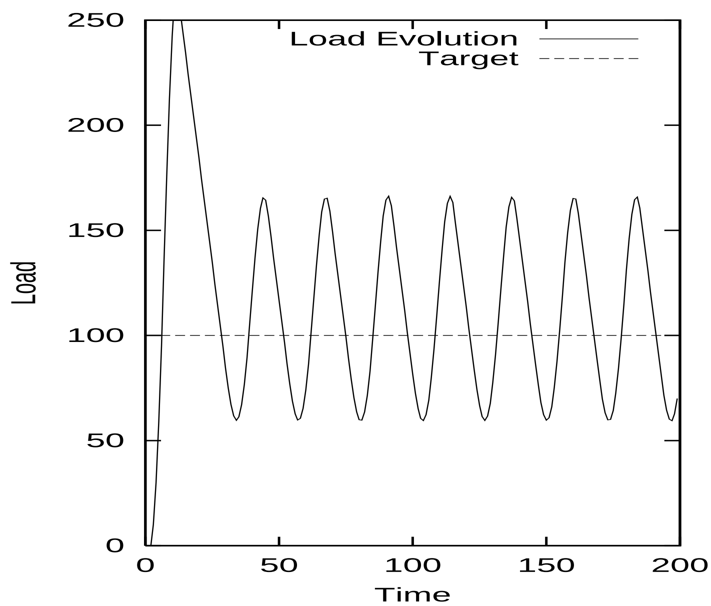

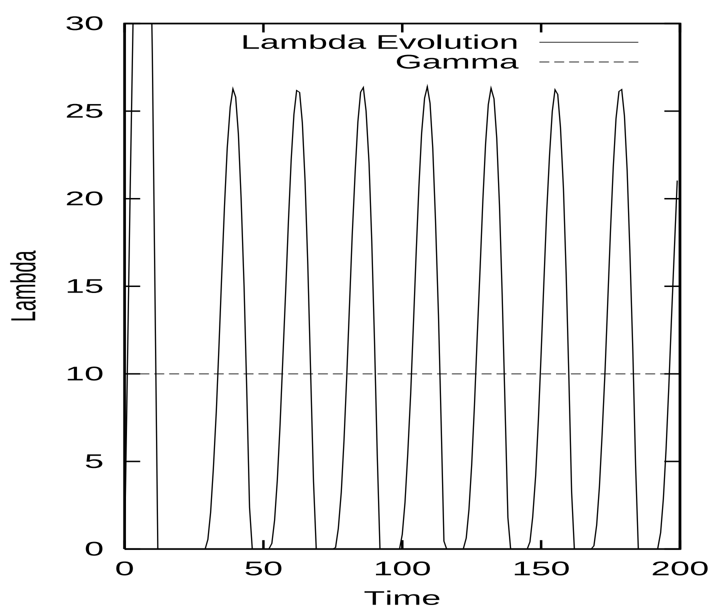

With $\varepsilon = 0.5$:

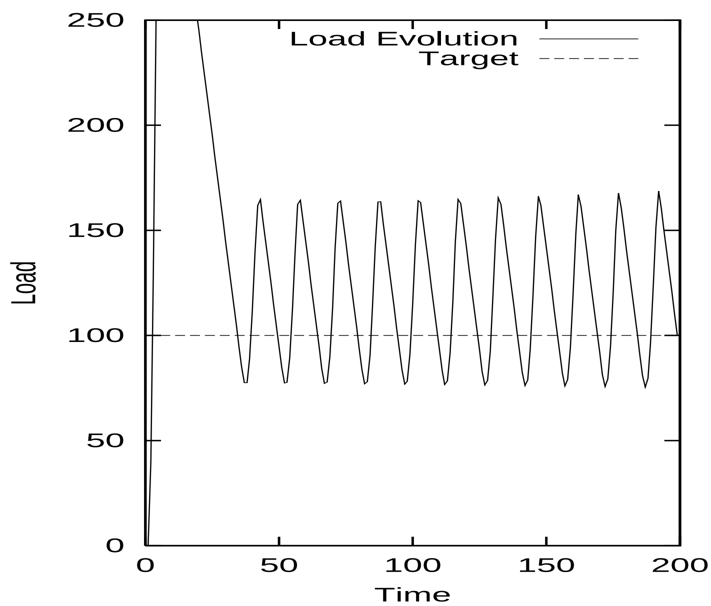

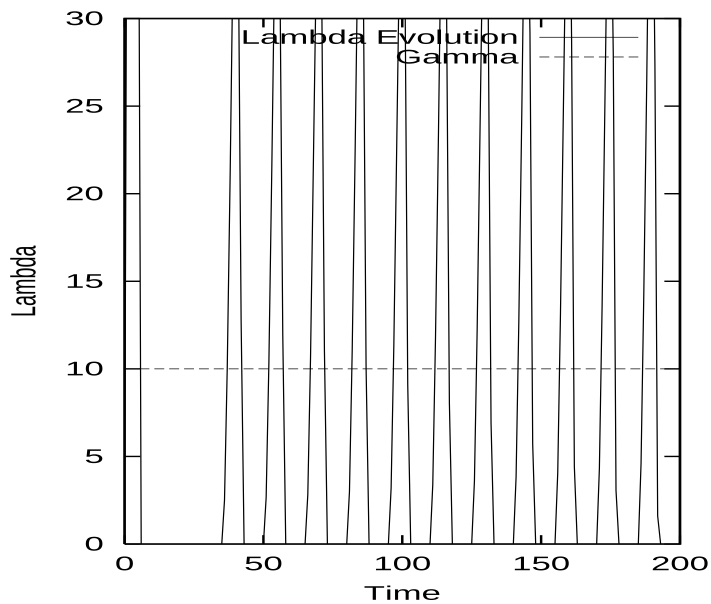

Answer: no

$\longrightarrow$ looks good but looks can be deceiving

Time to try something strange

$\longrightarrow$ any (crazy) ideas?

$\longrightarrow$ good for course project (assuming it works)

Odd looking modification to `Method C`:

`Method D`:

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

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

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

$\longrightarrow$ what's going on?

Sanity check: at desired operating point $Q(t) = Q^*$ and $\lambda(t) = \gamma$, nothing should move

$\longrightarrow$ check with methods A, B and C

$\longrightarrow$ fixed-point property

$\longrightarrow$ what about `Method D`?

Now: does `Method D` get to the targe fixed point?

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

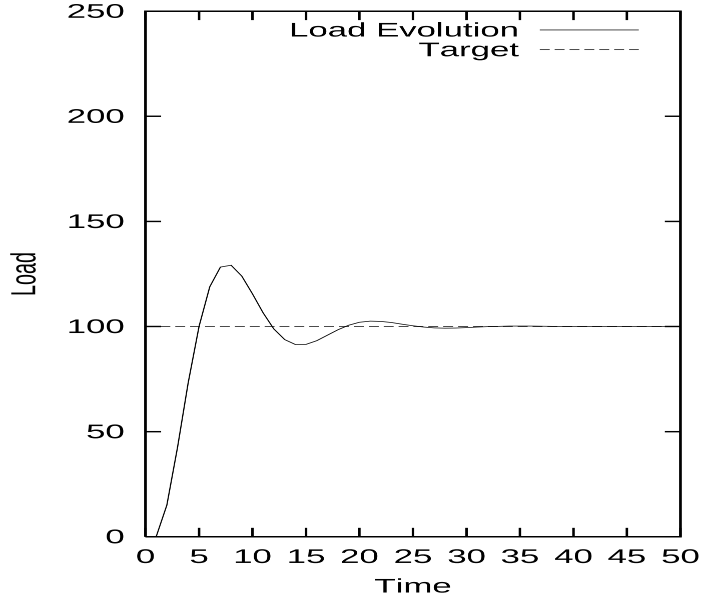

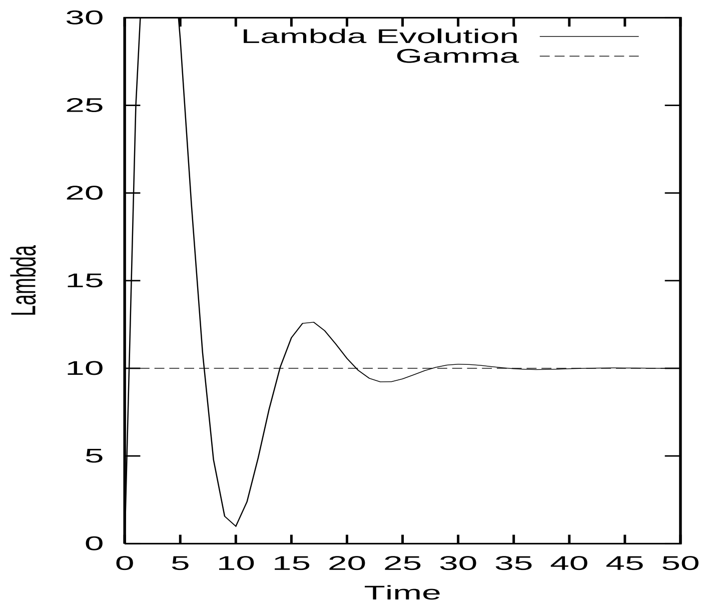

With $\varepsilon = 0.5$ and $\beta = 1.1$:

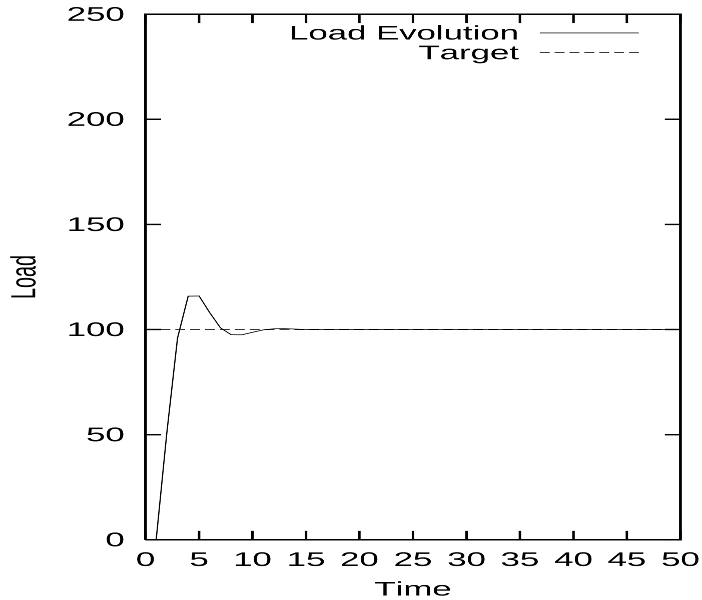

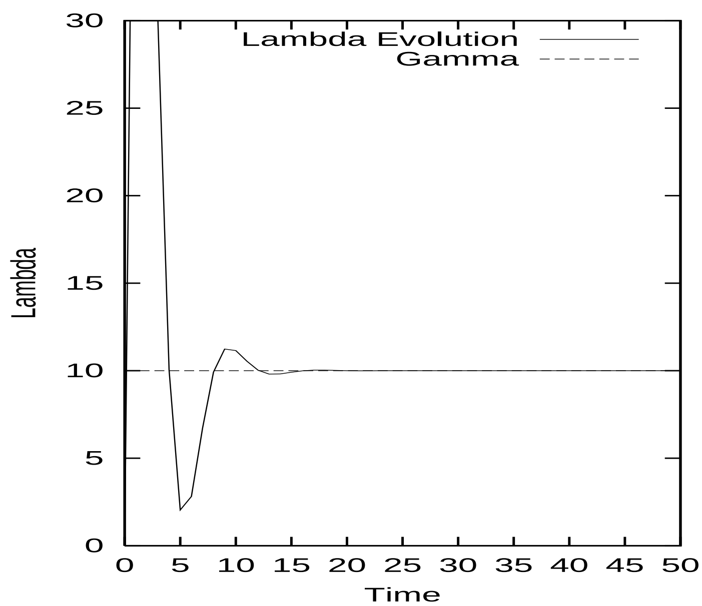

With $\varepsilon = 0.1$ and $\beta = 1.0$:

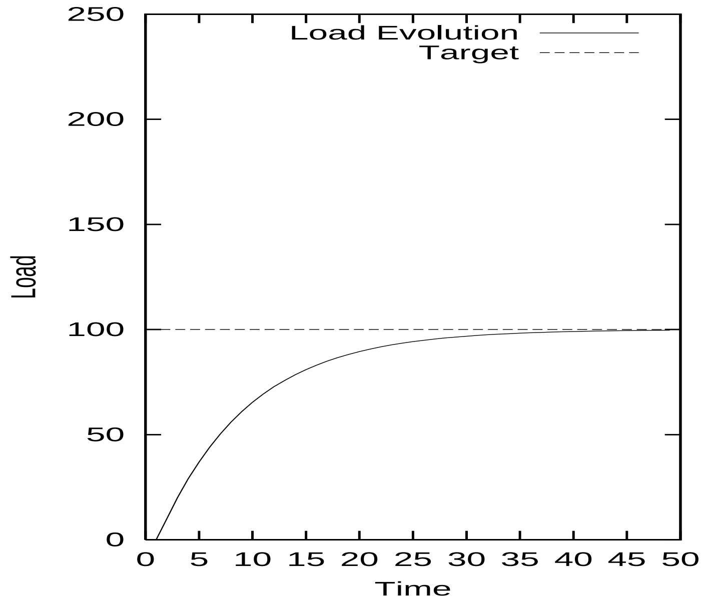

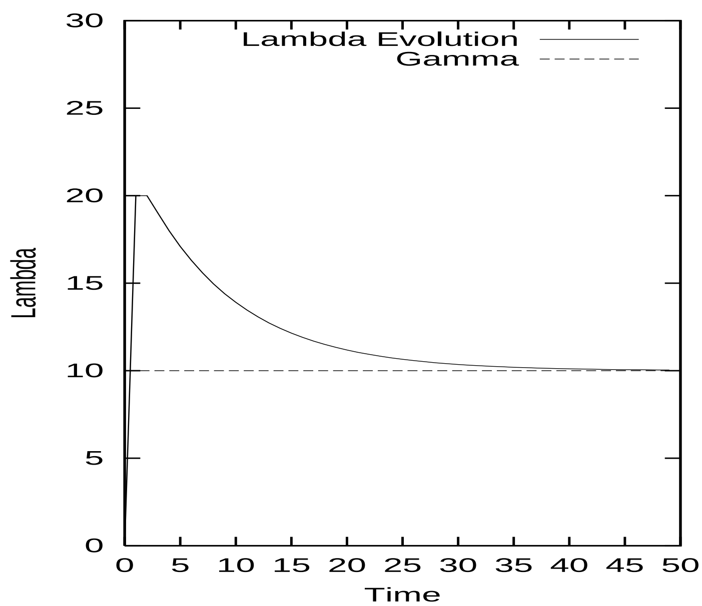

Remarks:

- Method D has desired behavior
- Is superior to Methods A, B, and C
- No unbounded oscillation
- In fact, dampening and convergence to desired operating point

  $\longrightarrow$ converges to target operating point (Q^*, γ)

  $$\lim_{t\to\infty}(Q(t), \lambda(t)) = (Q^*, \gamma)$$

  $\longrightarrow$ asymptotic stability
- Starting point $(Q(0), \lambda(0))$ issue:

  $\longrightarrow$ if target is reached from anywhere: global stability

  $\longrightarrow$ if target is reached when nearby: local stability

  $\longrightarrow$ want global stability

Why does it work?

What is the role of the $-\beta(\lambda(t) - \gamma)$ term in the control law:

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

Need to look beneath the hood ...

$\longrightarrow$ do you care about the engine or just the exterior?

$\longrightarrow$ are you "deep" or superficial?

$\longrightarrow$ answer: let's try to be both