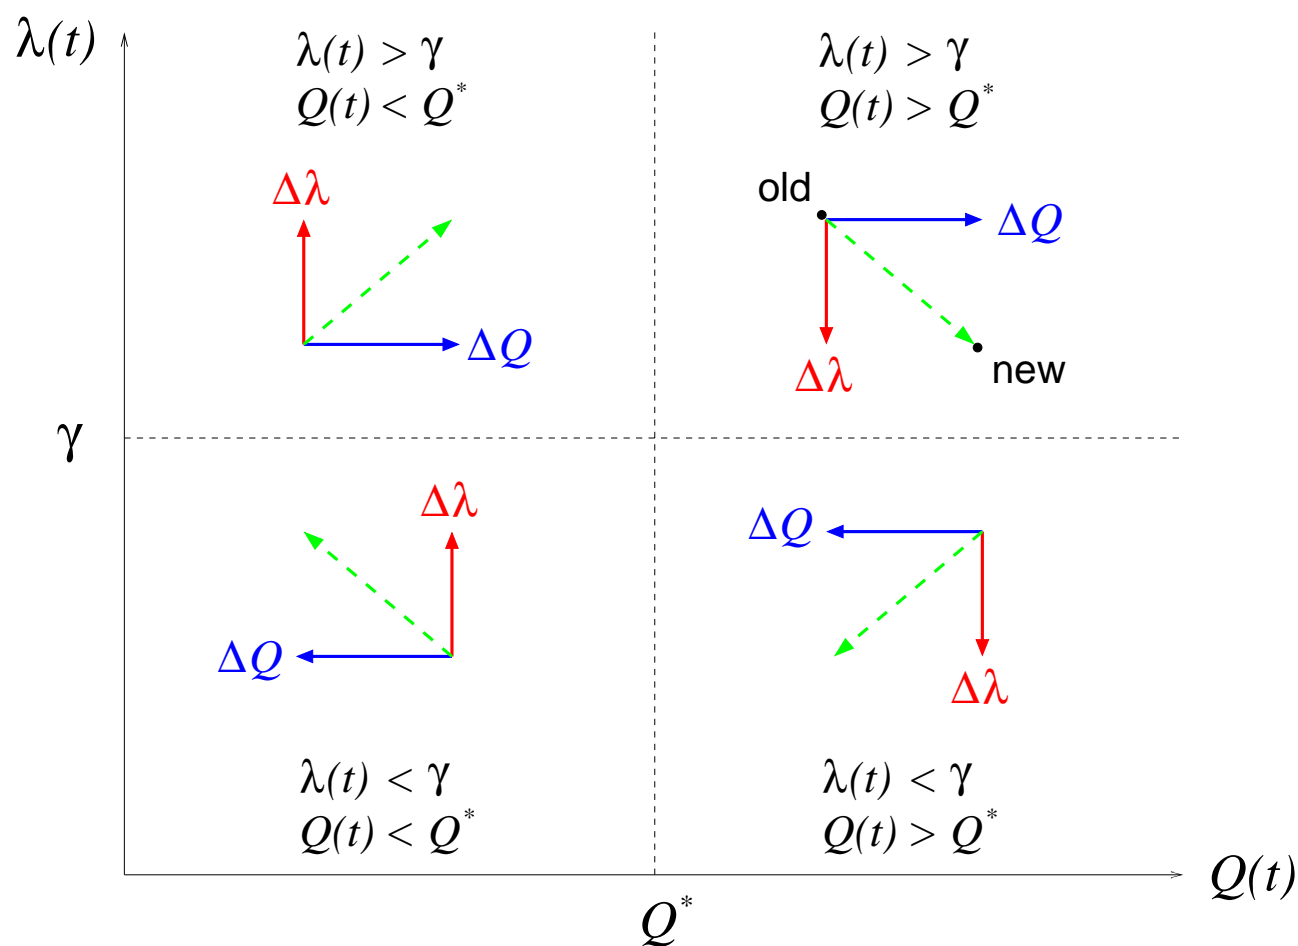



Visualize action in 2-D $(Q(t), \lambda(t))$ -space:

→ phase space




Convergent trajectory:

→ asymptotically stable & optimal




Divergent trajectory:

→ unstable



Stable (but not asymptotically so) trajectory:

→ limit cycle



Which case arises depends on the specifics of protocol actions.

For example:

- Methods A and C: divergent
- Method B: stable (but not asymptotically)
→ TCP
- Method D: asymptotically stable & optimal
→ “optimal control”

Why does Method D work:

→ overview of underlying mathematics

First, represent in continuous form:

→ easier to manipulate

→ more elegant

$$\begin{aligned}\frac{dQ(t)}{dt} &= \lambda(t) - \gamma \\ \frac{d\lambda(t)}{dt} &= \varepsilon(Q^* - Q(t)) - \beta(\lambda(t) - \gamma)\end{aligned}$$

In vector form:

$$\frac{d}{dt} \begin{pmatrix} Q(t) \\ \lambda(t) \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -\varepsilon & -\beta \end{pmatrix} \begin{pmatrix} Q(t) \\ \lambda(t) \end{pmatrix} + \begin{pmatrix} -\gamma \\ \varepsilon Q^* + \beta\gamma \end{pmatrix}$$

View as

$$\frac{d\mathbf{x}(t)}{dt} = A\mathbf{x}(t) + \mathbf{b}$$

where $\mathbf{x}(t) = (Q(t), \lambda(t))$.

Note: at desired/optimal operating point (Q^*, γ) ,

$$\frac{dQ(t)}{dt} = 0 \quad \text{and} \quad \frac{d\lambda(t)}{dt} = 0$$

→ (Q^*, γ) is a rest point or equilibrium

→ “when in heaven/nirvana, stay put”

Rest point (or fixed-point) condition is necessary but not sufficient:

→ external factors & noise makes \mathbf{x} stray

→ can it return to (Q^*, γ)

→ question of stability

Ex.: ball on a hill vs. ball in a valley



- both are fixed-points
- only one is stable w.r.t. perturbation

How do we determine stability of

$$d\mathbf{x}(t)/dt = A\mathbf{x}(t) + \mathbf{b} ?$$

Idea: consider $dz/dt = az$

- a : positive or negative constant
- what's the rest point?
- what happens to z over time?
- a : eigenvalue of the (1-D) system

Same idea applies to multi-dimensional linear systems

- in n dimensions: n eigenvalues
- asymptotic stability: all are negative
- stability: some may be 0
- unstable: one or more positive eigenvalues

What are eigenvalues, eigenvectors. . .

- given matrix A : $A\mathbf{u} = a\mathbf{u}$
- \mathbf{u} : eigenvector of A
- a : \mathbf{u} 's eigenvalue
- “operator” A stretches or pulls

Thus: for our 2-D congestion control system, stability depends on the eigenvalues of

$$A = \begin{pmatrix} 0 & 1 \\ -\varepsilon & -\beta \end{pmatrix}$$

Congestion control system is asymptotically stable if the real parts of the eigenvalues are strictly negative

→ eigenvalues can be complex

What remains... let's calculate and check!

Eigenvalues are obtained from characteristic equation

$$\begin{aligned}\det \begin{pmatrix} -\nu & 1 \\ -\varepsilon & -\beta - \nu \end{pmatrix} &= \nu(\beta + \nu) + \varepsilon \\ &= \nu^2 + \beta\nu + \varepsilon = 0\end{aligned}$$

This yields:

$$\nu = \frac{-\beta}{2} \pm \frac{\sqrt{\beta^2 - 4\varepsilon}}{2}$$

If $\varepsilon > 0$, then $Re(\nu) < 0$. If, in addition, $0 < \varepsilon < \beta/2$ then the eigenvalues are real.

Check: without the $-\beta(\lambda(t) - \gamma)$ term in the control law, the characteristic equation becomes

$$\nu^2 + \varepsilon = 0$$

Thus, $\nu = \pm\sqrt{\varepsilon}$ and the real part of ν is zero, violating the asymptotic stability condition.

→ mathematical requirement

→ intuitively: damping effect

What about when throughput γ experiences congestion: under excessive load goes down

→ note: we assumed constant γ in analysis

TCP congestion control

Recall:

$$\text{EffectiveWindow} = \text{MaxWindow} - (\text{LastByteSent} - \text{LastByteAcked})$$

where

$$\text{MaxWindow} = \min\{ \text{AdvertisedWindow}, \text{CongestionWindow} \}$$

Key question: how to set `CongestionWindow` which, in turn, affects ARQ's sending rate?

- linear increase/exponential decrease
- AIMD

TCP congestion control components:

(i) Congestion avoidance

→ linear increase/exponential decrease

→ additive increase/exponential decrease (AIMD)

As in Method B, increase `CongestionWindow` linearly,
but decrease exponentially

Upon receiving ACK:


$$\text{CongestionWindow} \leftarrow \text{CongestionWindow} + 1$$

Upon timeout:

$$\text{CongestionWindow} \leftarrow \text{CongestionWindow} / 2$$


But is it correct...

“Linear increase” time diagram:



→ results in exponential increase

What we want:



→ increase by 1 every window

Thus, linear increase update:

$$\begin{aligned} \text{CongestionWindow} &\leftarrow \text{CongestionWindow} \\ &\quad + (1 / \text{CongestionWindow}) \end{aligned}$$

Upon timeout and exponential backoff,

$$\text{SlowStartThreshold} \leftarrow \text{CongestionWindow} / 2$$

(ii) Slow Start

Reset `CongestionWindow` to 1

Perform exponential increase

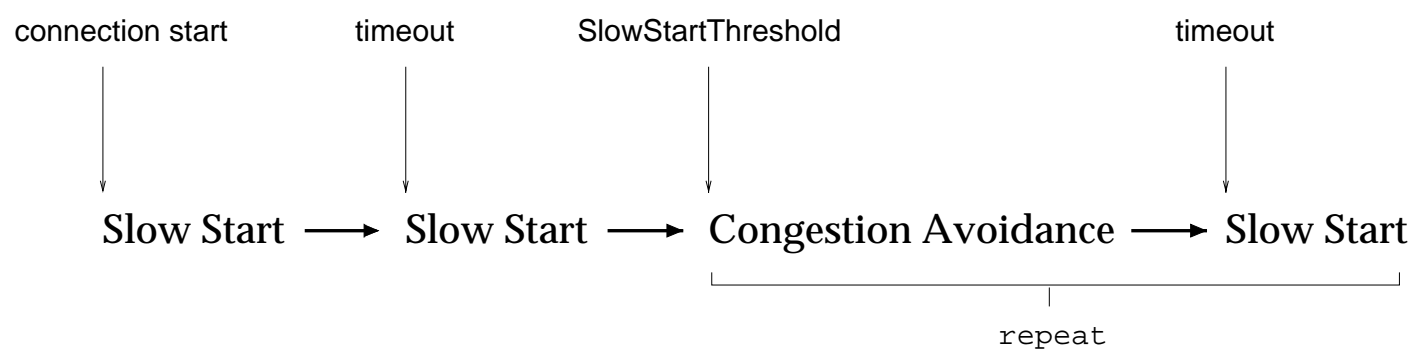
$\text{CongestionWindow} \leftarrow \text{CongestionWindow} + 1$

- Until timeout at start of connection
 - rapidly probe for available bandwidth
- Until `CongestionWindow` hits `SlowStartThreshold` following Congestion Avoidance
 - rapidly climb to safe level
 - “slow” is a misnomer
 - exponential increase is super-fast

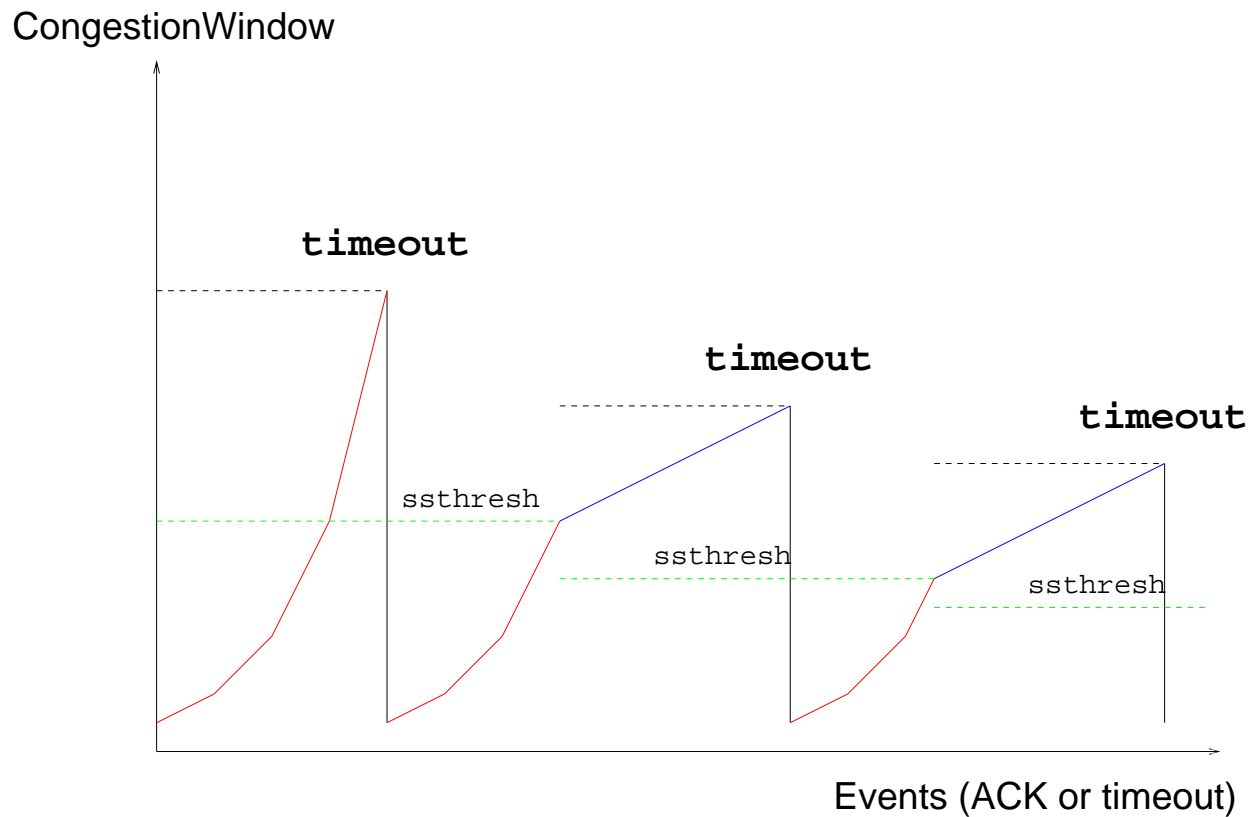
Basic dynamics:

→ after connection set-up

→ before connection tear-down



CongestionWindow evolution:



- what happens if receiver window size hits max?
- DOE, supercomputing centers, etc.

(iii) Exponential timer backoff

$\text{TimeOut} \leftarrow 2 \cdot \text{TimeOut}$ if retransmit

(iv) Fast Retransmit

Upon receiving three duplicate ACKs:

- Transmit next expected segment
 - segment indicated by ACK value
- Perform exponential backoff and commence Slow Start
 - three duplicate ACKs: likely segment is lost
 - react before timeout occurs

TCP Tahoe: features (i)-(iv)

(v) Fast Recovery

Upon Fast Retransmit:

- Skip Slow Start and commence Congestion Avoidance
→ dup ACKs: likely spurious loss
- Insert “inflationary” phase just before Congestion Avoidance

Inflationary phase:

- $\text{SlowStartThreshold} \leftarrow \text{CongestionWindow} / 2$
- $\text{CongestionWindow} \leftarrow \text{SlowStartThreshold} + 3$
- On each additional duplicate ACK, increment CongestionWindow
- On first non-dup ACK, commence Congestion Avoidance
 $\text{CongestionWindow} \leftarrow \text{SlowStartThreshold}$

TCP Reno: features (i)-(v)

→ pre-dominant form

Many more versions of TCP:

→ NewReno w/ SACK, w/o SACK, Vegas, etc.

→ wireless, ECN, multiple time scale

→ mixed verdict; pros/cons

Given sawtooth behavior of TCP's linear increase/exponential backoff:

Why use exponential backoff and not Method D?

- For multimedia streaming (e.g., pseudo real-time), AIMD (Method B) is not appropriate
→ use Method D
- For unimodal case—throughput decreases when system load is excessive—story is more complicated
→ asymmetry in control law needed for stability