Transport

Review: slow start, additive inc

AI is slow.
Persistent connections
Large window sizes
Different laws to evolve congestion window

Say `MSS` = 1 KByte
Default `ssthresh` = 64KB = 64 `MSS`

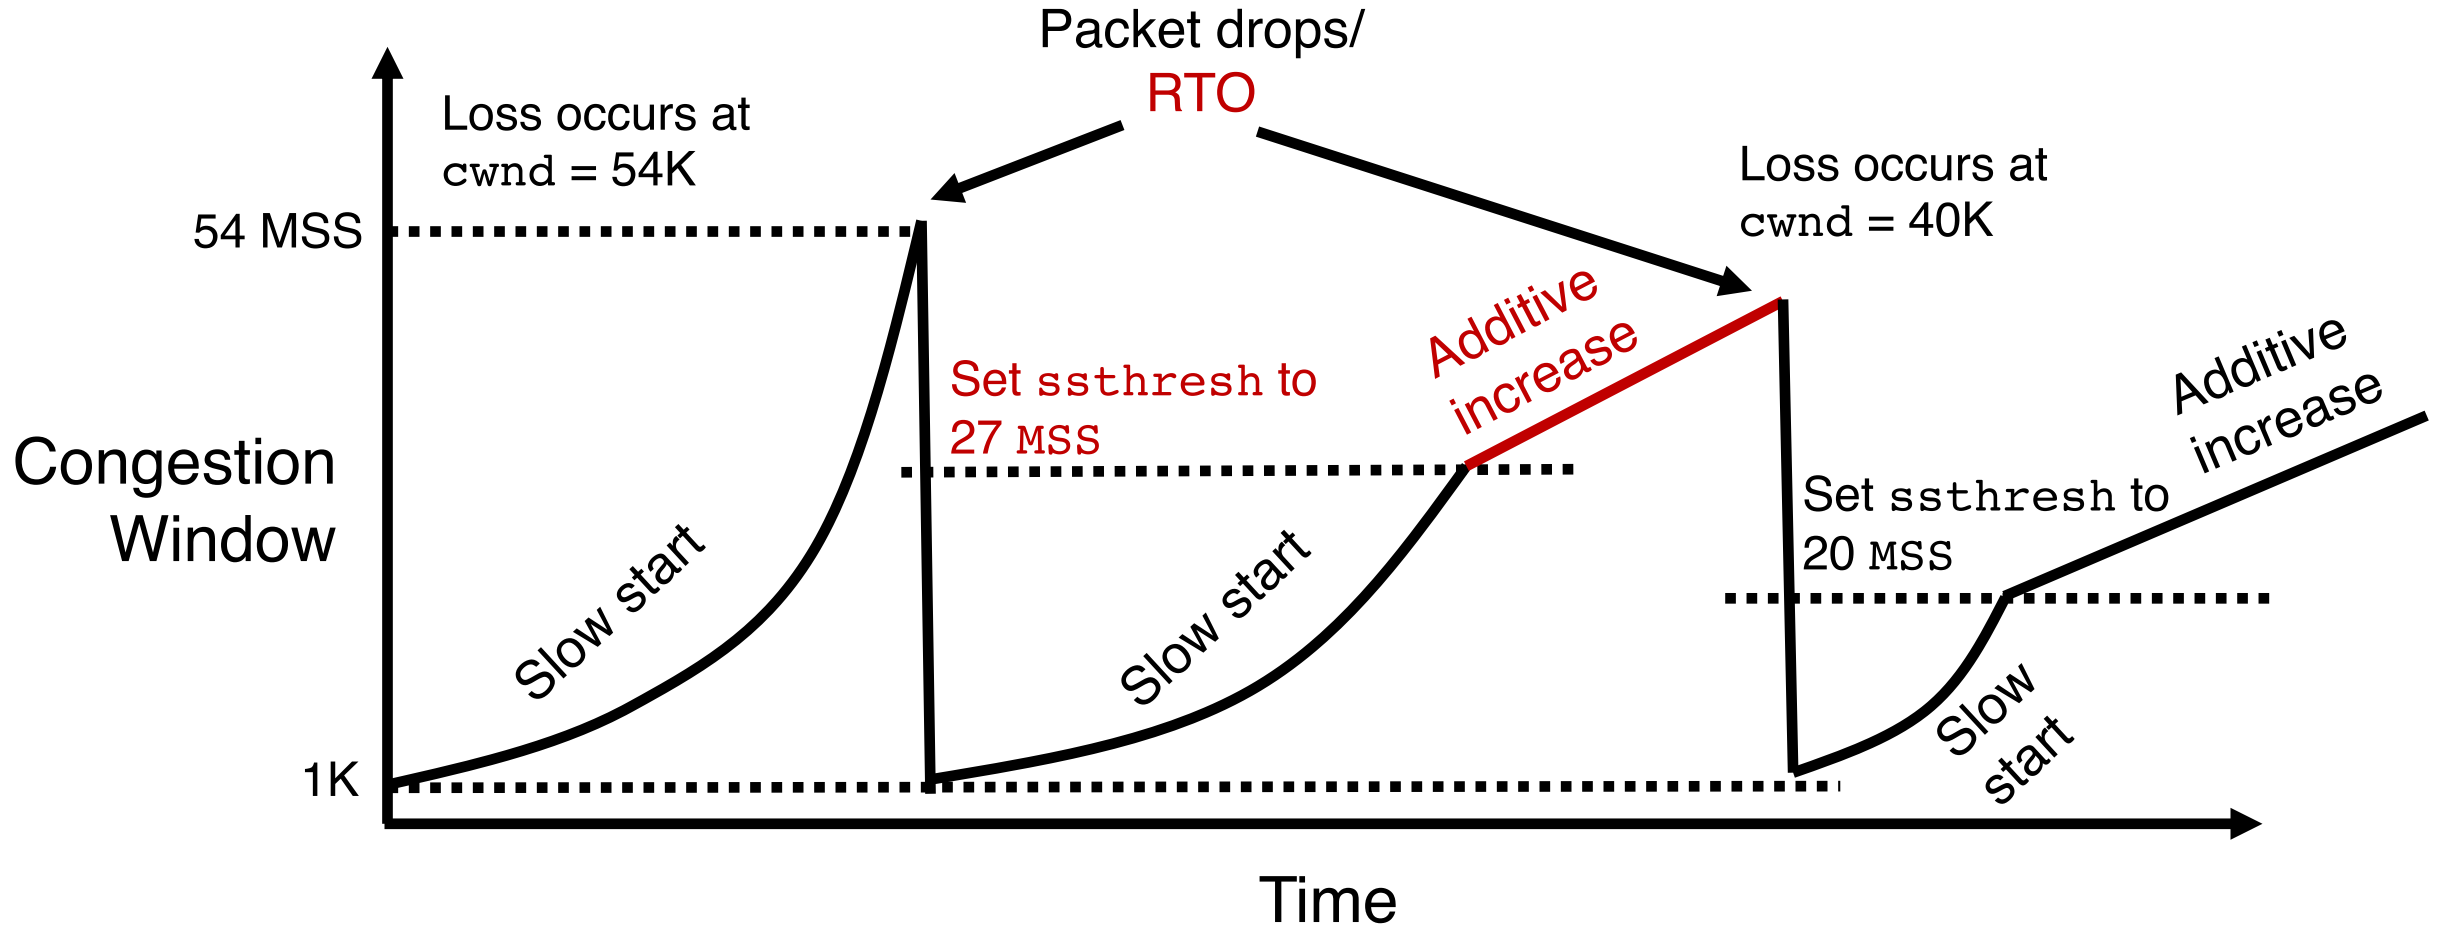

The components of delay

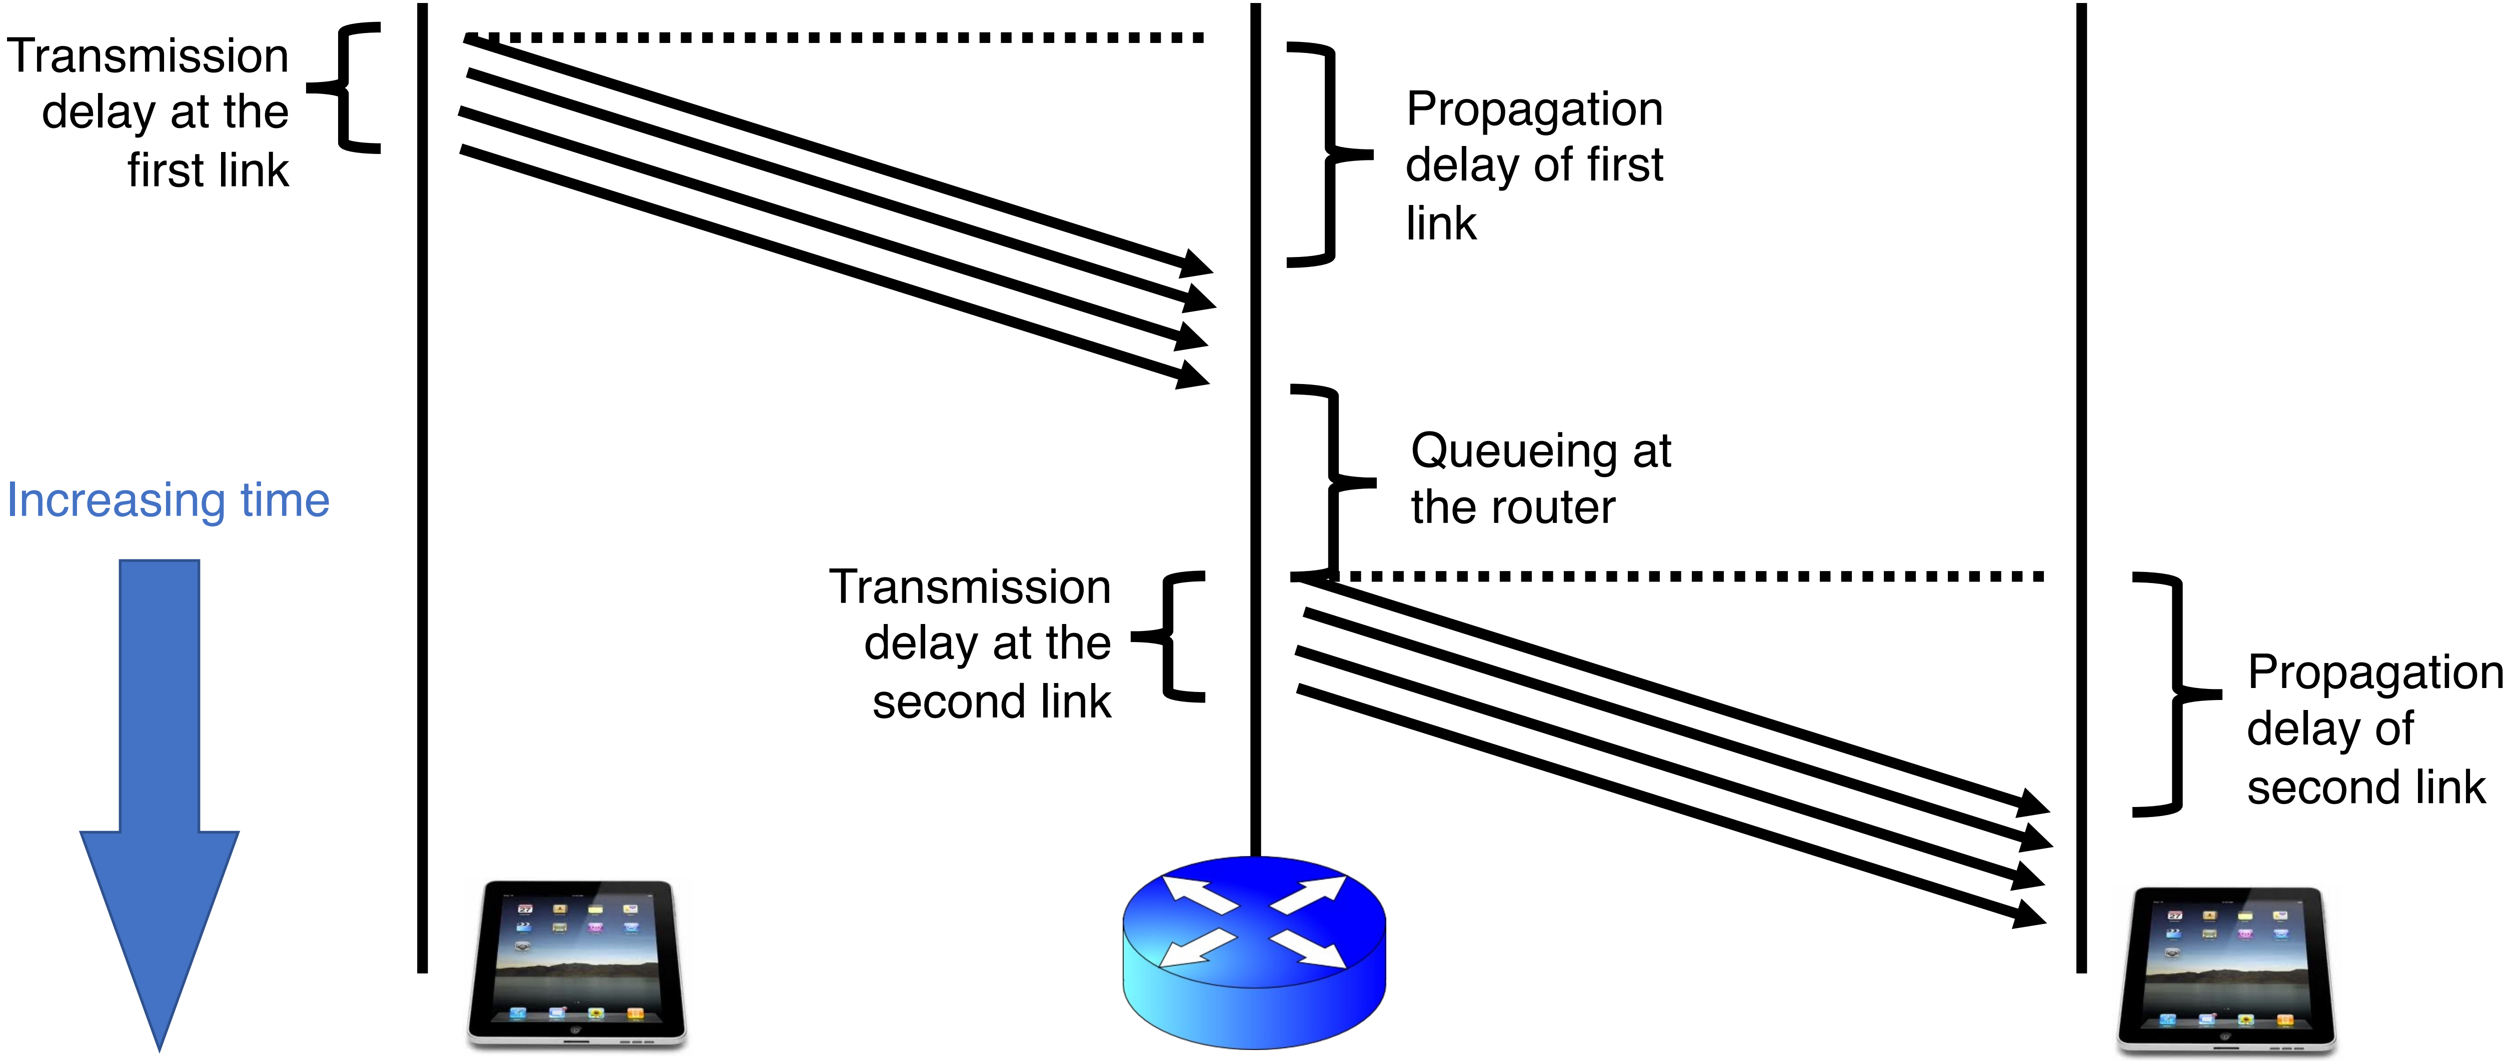

Bandwidth-Delay Product

Steady state `cwnd` for a single flow

- Suppose the bottleneck link has rate C
- Suppose the propagation round-trip delay (propRTT) between sender and receiver is T
- Ignore transmission delays for this example;
- Assume steady state: highest sending rate with no bottleneck congestion
- Q: how much data is in flight over a single RTT?
- C * T data i.e., amount of data unACKed at any point in time
- ACKs take time T to arrive (without any queueing). In the meantime, sender is transmitting at rate C

The Bandwidth-Delay Product

- C * T = bandwidth-delay product:
  - The amount of data in flight for a sender transmitting at the ideal rate during the ideal round-trip delay of a packet

- Note: this is just the amount of data "on the pipe"

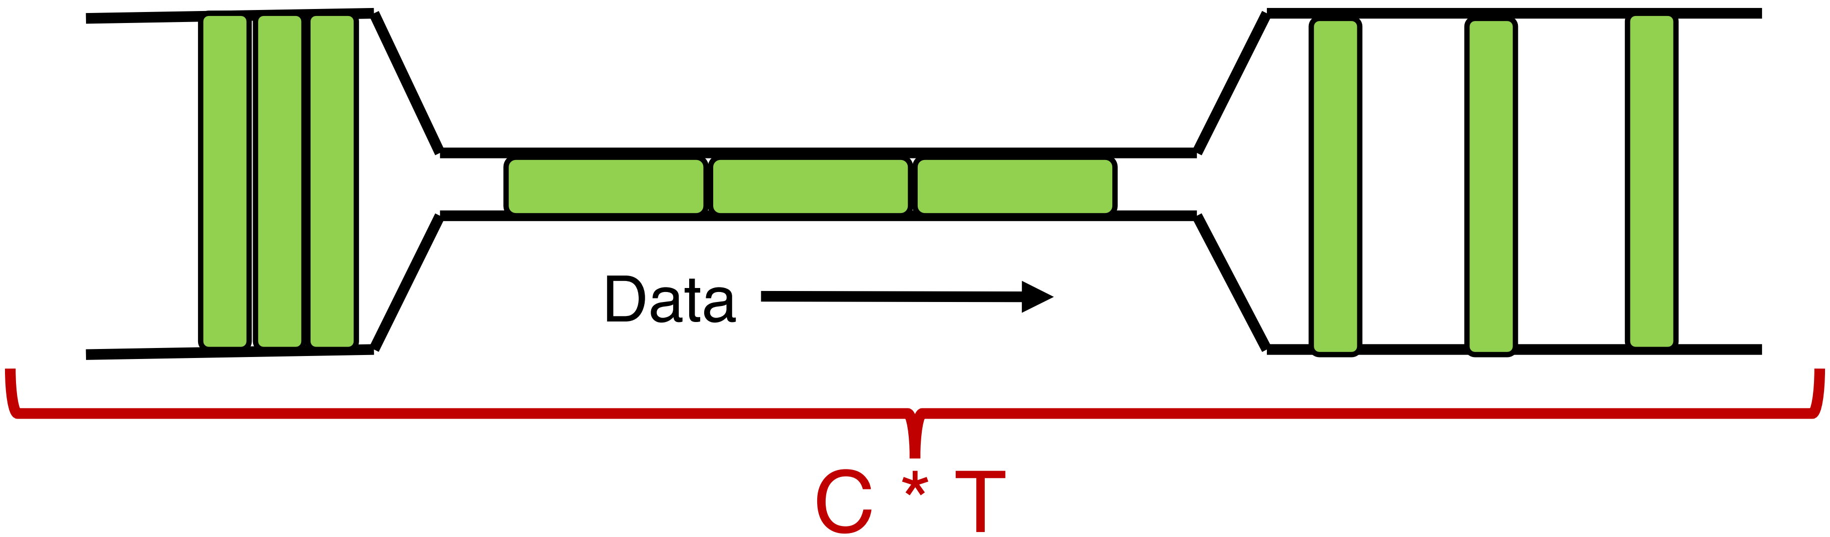

The Bandwidth-Delay Product

- Q: What happens if cwnd > C * T?
  - i.e., where are the rest of the in-flight packets?

- A: Waiting at the bottleneck router queues

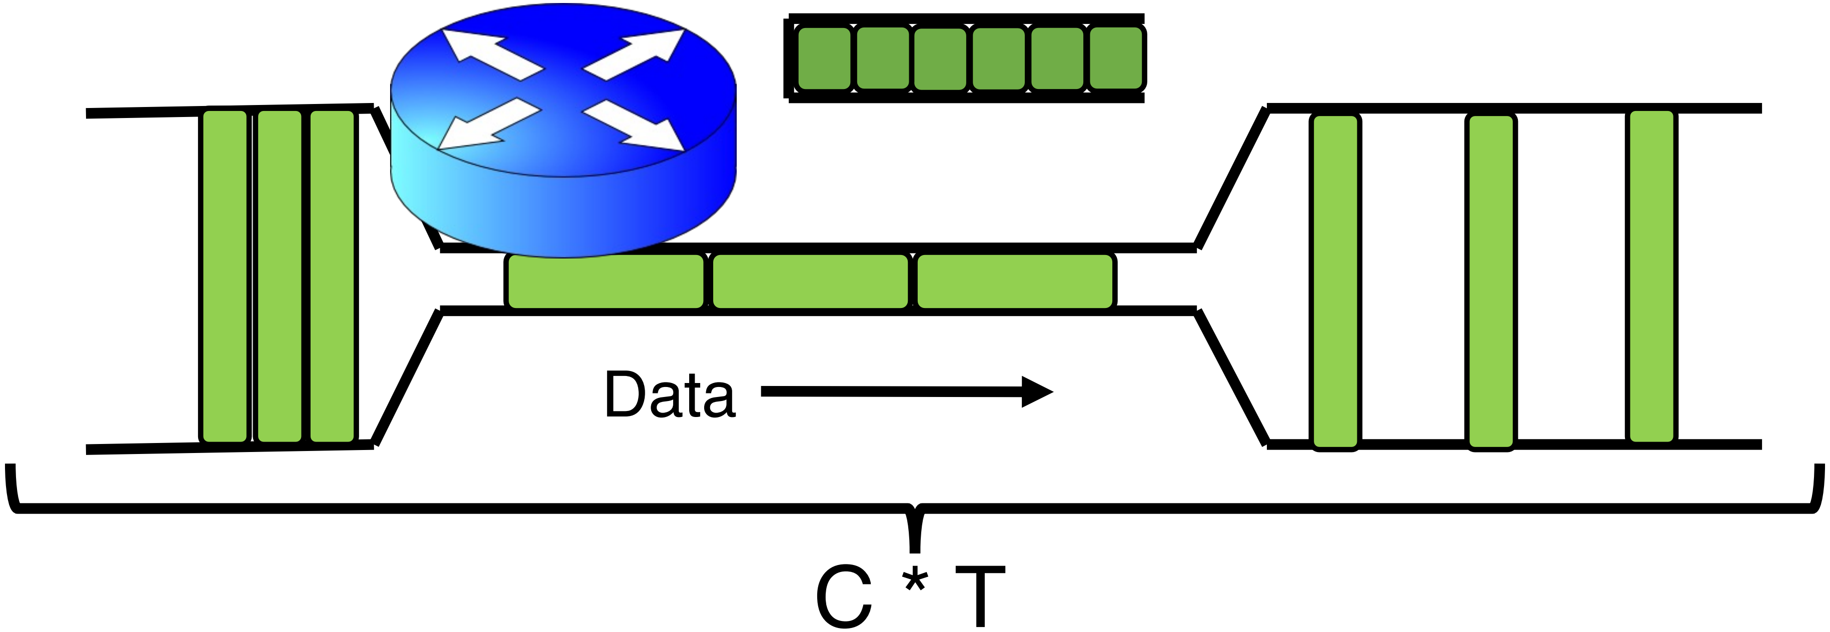

Router buffers and the max `cwnd`

- Router buffer memory is finite: queues can only be so long
  - If the router buffer size is B, there is at most B data waiting in the queue
- If cwnd increases beyond C * T + B, data is dropped!

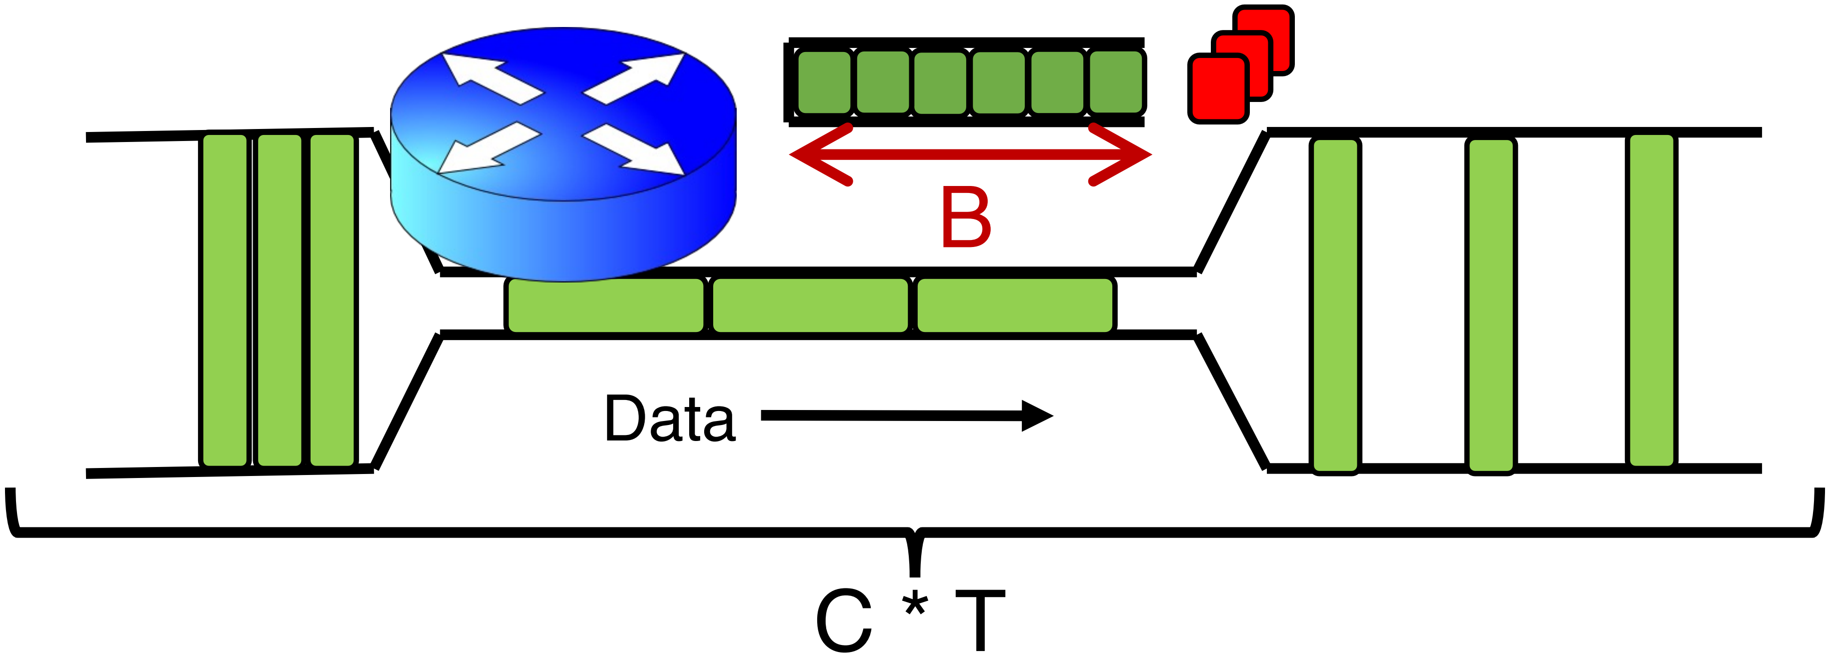

BDP is a crucial value for a flow

- Bandwidth-Delay Product (BDP) governs the window size of a single flow at steady state
- The bottleneck router buffer size governs how much the `cwnd` can exceed the BDP before packet drops occur
- BDP is the ideal desired window size to use the full bottleneck link, without any queueing.
  - Accommodating flow control, also the min socket buffer size to use the bottleneck link fully

Demo of the impact of BDP & B

- Utilization
- Congestion window

Detecting and Reacting Better to Packet Loss

Can we detect loss earlier than RTO?

- Key idea: use the information in the ACKs. How?
- Suppose successive (cumulative) ACKs contain the same ACK#
  - Also called duplicate ACKs
  - Occur when network is reordering packets, or one (but not most) packets in the window were lost
  - Fast retransmit: (1) Immediately retransmit packet
- Reduce `cwnd` when you see many duplicate ACKs
  - Consider many dup ACKs a strong indication that packet was lost
  - Default threshold: 3 dup ACKs, i.e., triple duplicate ACK
  - Make cwnd reduction gentler than setting cwnd = 1; recover faster
  - Fast retransmit: (2) reduce window to half of its current value

Additive Increase/Multiplicative Decrease

Say `MSS` = 1 KByte

Default `ssthresh` = 64KB = 64 `MSS`

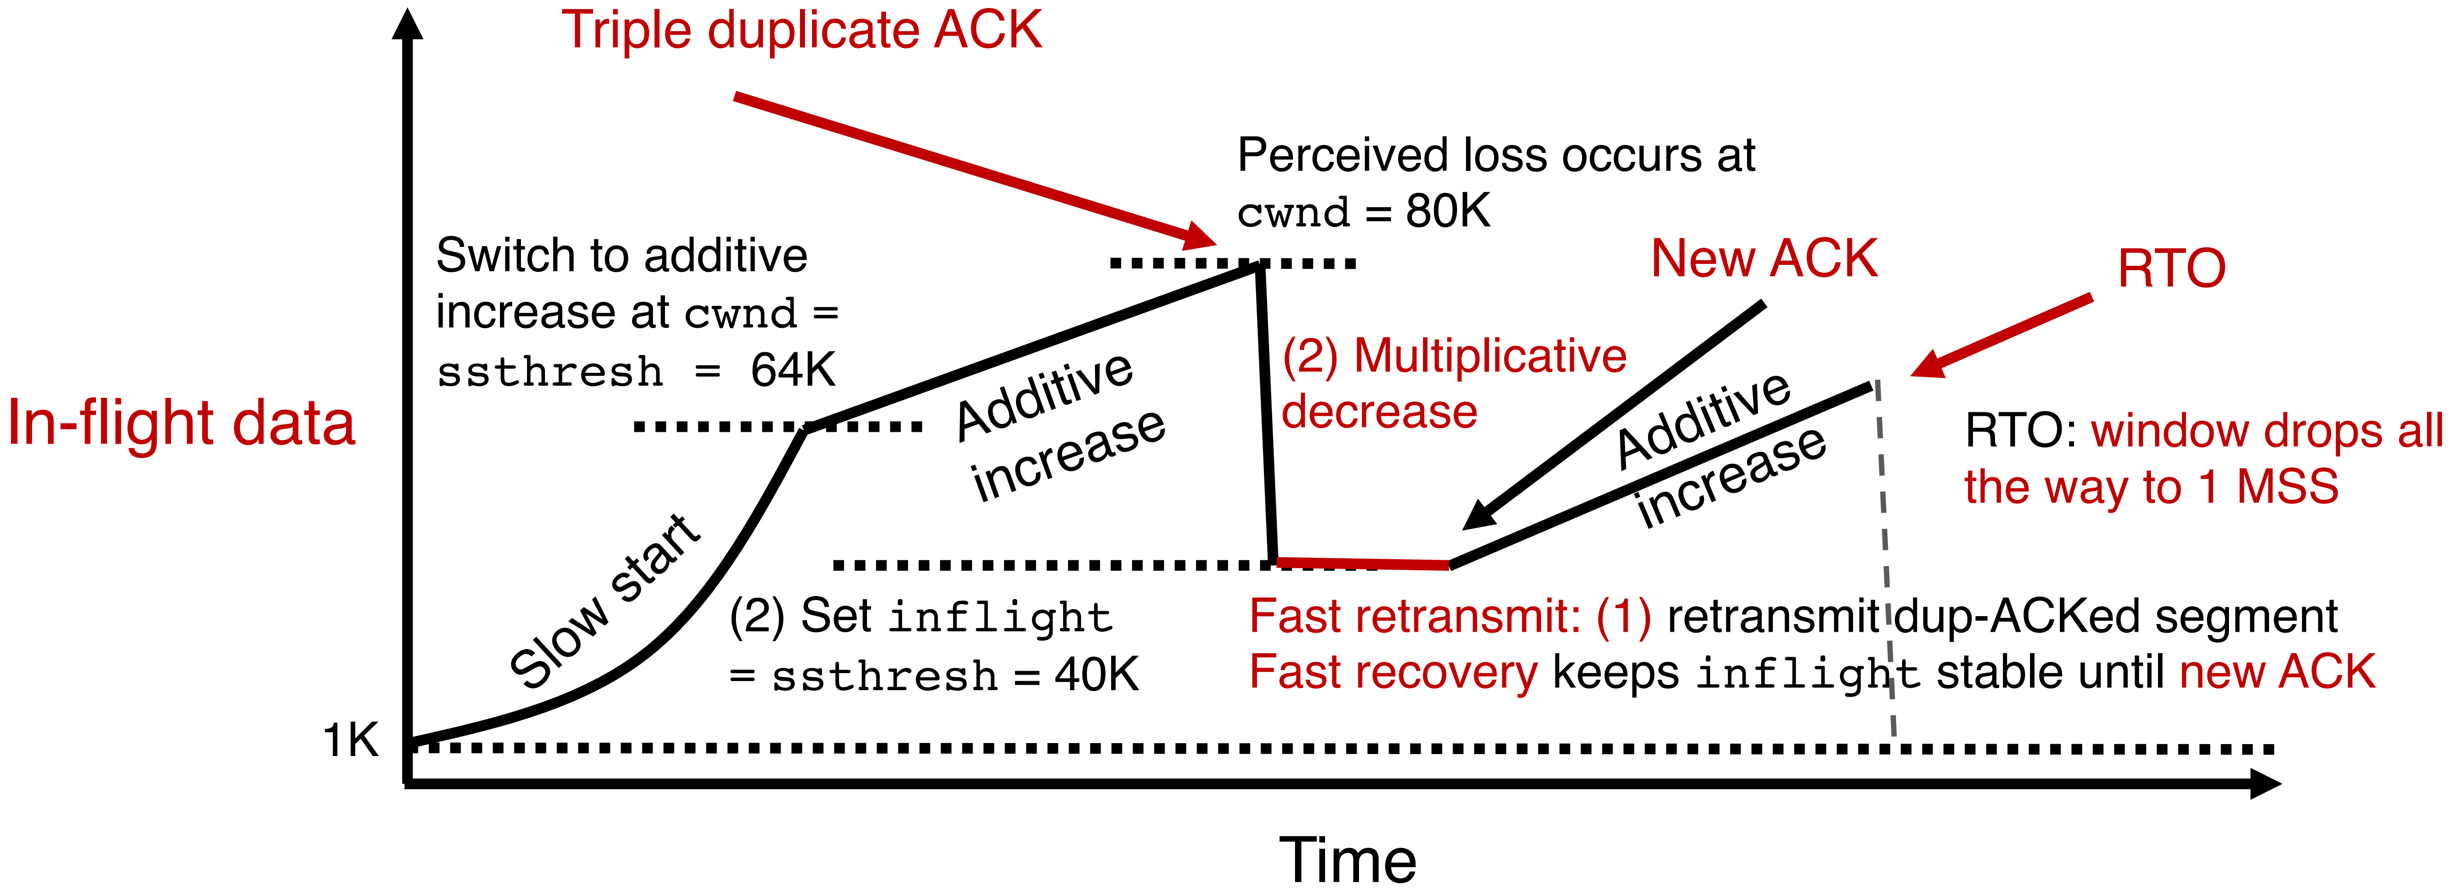

TCP New Reno performs additive increase and multiplicative decrease of congestion window.

In short, we often refer to this as AIMD.

Multiplicative decrease is a part of all TCP algorithms. It is necessary for fairness across TCP flows.

Why does multiplicative decrease help?

Efficiency and Fairness

Chiu and Jain, “Increase and decrease algorithms for congestion avoidance”

Efficient allocation

- Don't want sources to transmit either too slow or too fast
  - Slow: Underutilize the network
  - Fast: High delays, lose packets
- Every endpoint is reacting
  - May *all* under/overshoot
  - Large oscillations possible!
- Optimal efficiency:
  - $\Sigma x_i = X_{goal}$ e.g., link capacity
- Efficiency = 1 - distance from efficiency line

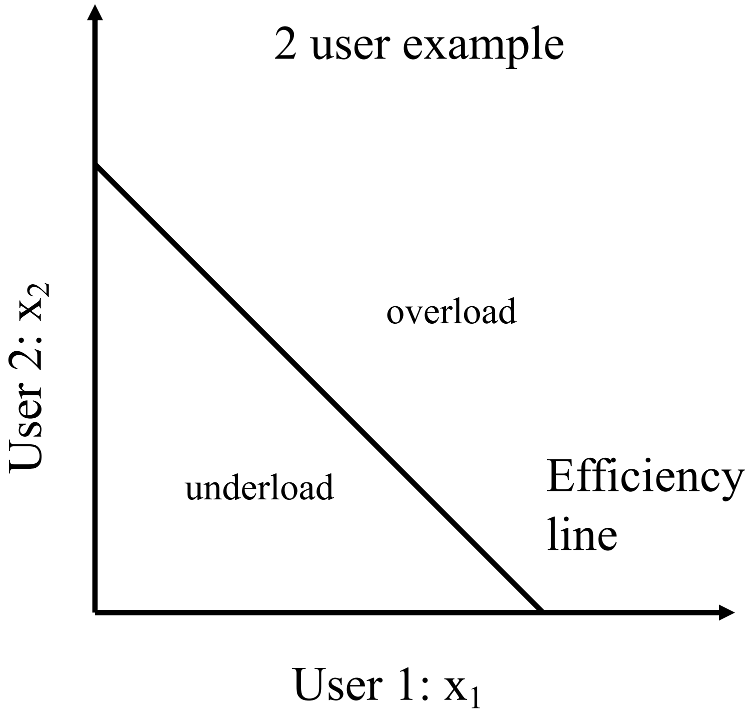

Fair allocation

- Max-min fairness
- Flows which share the same bottleneck get the same amount of bandwidth

$$F(x) = \frac{\left(\sum x_i\right)^2}{n\left(\sum x_i^2\right)}$$

- Fairness = 1 - distance from fairness line

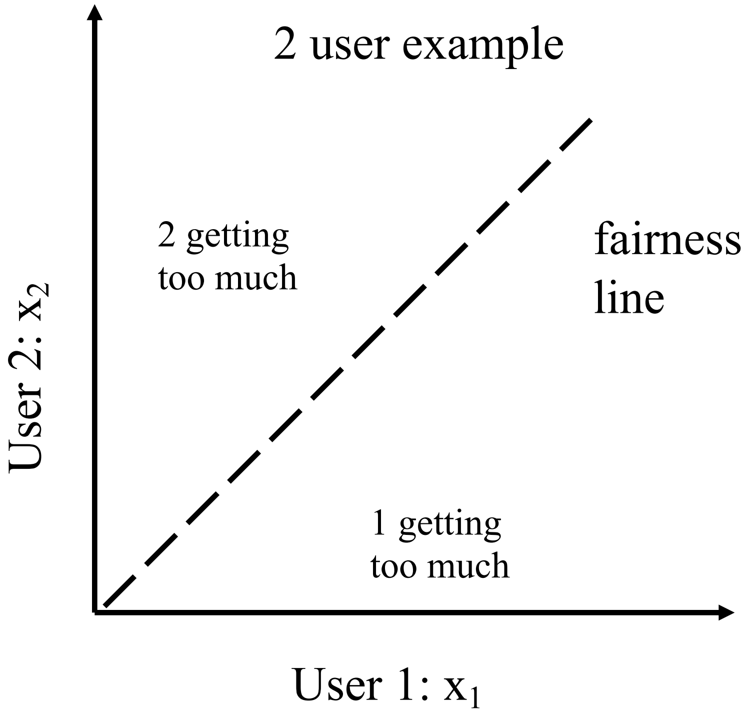

How should transports react?

- Given efficiency and fairness goals above, how should transports behave?

- Consider $x(t)$, window or rate of a source, evolving over time t

- Assume discrete time steps.
- $x(t + 1) =$ function of $x(t)$, feedback from the network

Linear control rules

$$x_i(t+1) = \begin{cases} a_I + b_I x_i(t) & increase \\ a_D + b_D x_i(t) & decrease \end{cases}$$

- $x_i(t)$: window or rate of the i^{th} user at time t
- a_I, a_D, b_I, b_D: constant increase/decrease coefficients
- Assumption: All users receive same network feedback
    - *Binary* feedback: sense congestion or available capacity
- Assumption: All users increase or decrease simultaneously

Additive increase, multiplicative decrease

- $b_I = 1$, $a_D = 0$
- Multiplicative decrease enables converging to fairness
- Oscillates around the most efficient point

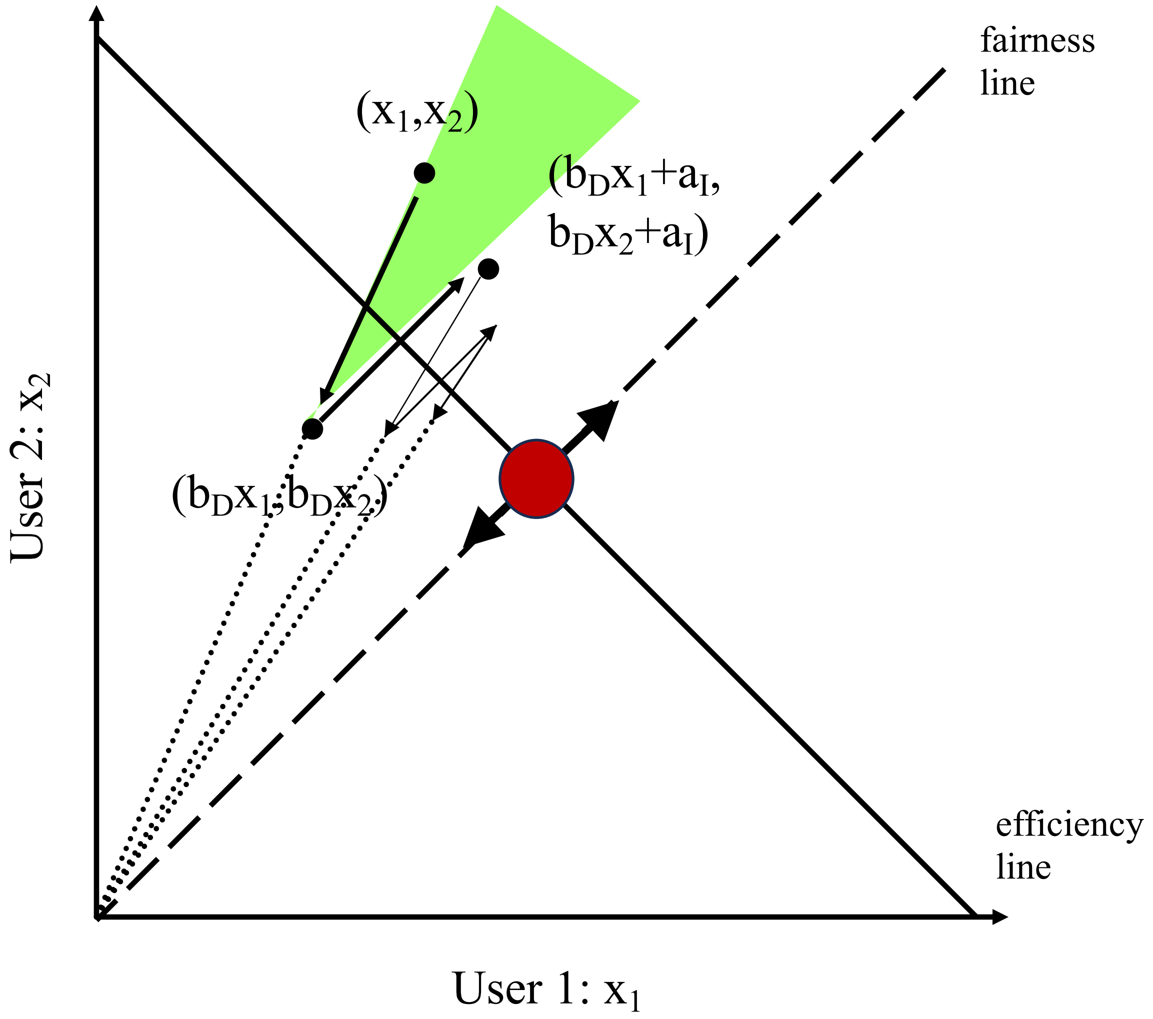

Convergent doesn't mean static