

More on Congestion


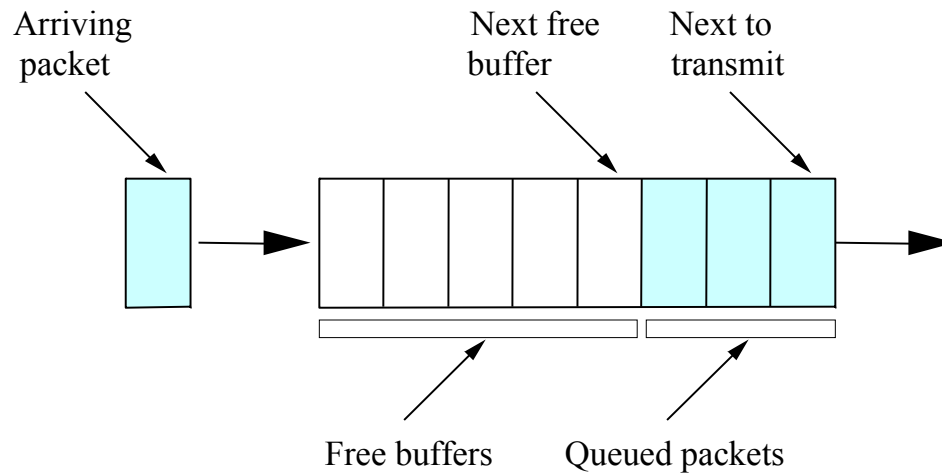
- Focus
 - How to we avoid congestion?
- Topics
 - Random Early Detection (RED) gateways
 - Explicit Congestion Notification (ECN)




Application
Presentation
Session
Transport
Network
Data Link
Physical

Congestion Control vs Avoidance

- TCP causes congestion as it probes for the available bandwidth and then recovers from it after the fact
 - Leads to loss, delay and bandwidth fluctuations
 - We want congestion avoidance, not congestion control
- Congestion avoidance mechanisms
 - Aim to detect incipient congestion, before loss.
 - Monitor queues to see that they absorb bursts, but not build steadily


Router Model: "FIFO with Tail Drop"






Incipient Congestion at a Router

- Sustained overload causes queue to build and overflow



Alternative: Random Early Detection (RED)

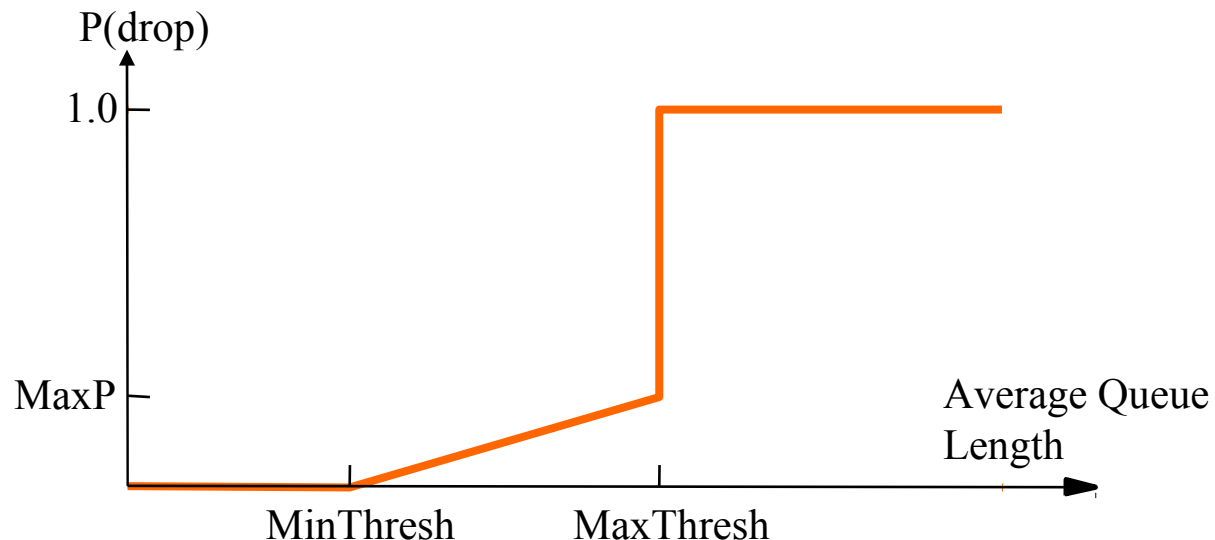
- Have routers monitor average queue and send “early” signal to source when it builds
 - _ probabilistically drop a packet



- Paradox: early loss can improve performance!

Red Drop Curve

- Start dropping a fraction of the traffic as queue builds
 - Expected drops proportional to bandwidth usage
 - When queue is too high, revert to drop tail
 - Nice theory, difficult to set parameters in practice



Explicit Congestion Notification (ECN)

- Why drop packets to signal congestion?
 - _ Drops are a robust signal, but there are other means ...
 - _ We need to be careful though: no extra packets
- ECN signals congestion with a bit in the IP header
- Receiver returns indication to the sender
 - _ Need to signal this reliably or we risk instability
- Sender reacts as though a loss has been detected

- RED actually works by “marking” packets
 - _ Mark can be a drop or ECN signal if hosts understand ECN
 - _ Supports congestion avoidance without loss

Aside: TCP Vegas (Peterson '94)

- RED needs router upgrades but no host upgrades
- Instead, can we upgrade host but not router?
- TCP Vegas looks at the difference between cwnd (the amount of outstanding data in the network) and that acknowledged from the other side in the last interval
 - Excess must be buffered in the network at router queues
 - Vegas slows down when it believes there is a queue and otherwise increases to use the available bandwidth
 - Requires accurate estimation of uncongested performance