

Congestion Control

Lecture 17

<http://www.cs.rutgers.edu/~sn624/352-F24>



Srinivas Narayana

= window size


Proportional to **throughput**

How much data to keep in flight?

Stop and Wait



Pipelined Reliability



It is difficult to know where the **bottleneck** link is

It is difficult to know how many other endpoints are using that link

Endpoints may join and leave at any time

Network paths may change over time, leading to different bottleneck links (with different link rates) over time

How to share a network,
with so many unknowns?

The approach that the Internet takes is to use a distributed algorithm to converge to an efficient and fair outcome.

The approach that the Internet takes is to use a **distributed algorithm** to converge to an efficient and fair outcome.

No one can centrally view or control all the endpoints and bottlenecks in the Internet.

Every endpoint must try to reach a globally good outcome by itself: i.e., in a distributed fashion.

This also puts a lot of **trust in endpoints**.

The approach that the Internet takes is to use a distributed algorithm to converge to an **efficient** and fair outcome.

If there is spare capacity in the bottleneck link, the endpoints should use it.

The approach that the Internet takes is to use a distributed algorithm to converge to an efficient and **fair** outcome.

If there are N endpoints sharing a bottleneck link, they should be able to get **equitable** shares of the link's capacity.

For example: $1/N$ 'th of the link capacity.

Flow Control vs. Congestion Control

- Avoid overwhelming the
receiving application
- Sender is managing the
receiver's socket buffer

- Avoid overwhelming the
bottleneck network link
- Sender is managing the
bottleneck link capacity and
bottleneck router buffers

The approach that the Internet takes is to use a distributed algorithm to converge to an efficient and fair outcome.

How to achieve this?

Approach: sense and react

Example: showering: Want “just right” water

Use a **feedback loop** with signals and knobs



Signals and Knobs in Congestion Control

- **Signals**

- Packets being ACK'ed
- Packets being dropped (e.g. RTO fires)
- Packets being delayed (RTT)
- Rate of incoming ACKs

} **Implicit** feedback signals measured directly at sender.
(There are also explicit signals that the network might provide.)

- **Knobs**

- What can you change to “probe” the available bottleneck capacity?
 - **Window size**
- Suppose the receiver socket buffer size is unbounded
 - **Congestion window**: window size used for congestion control
- Increase window/sending rate: e.g., add x or multiply by a factor of x
- Decrease window/sending rate: e.g., subtract x or reduce by a factor of x

Sense and react, sure...but how?

- Where do you want to be?
 - The **steady state**
- How do you get there?
 - Congestion control algorithms
- Sense accurately & react accordingly




The Steady State

Efficiency for a single TCP connection

What does **efficiency** look like?

- Suppose we want to achieve an **efficient** outcome for **one** TCP connection by observing network signals from the endpoint




- Q: How should the endpoint behave **at steady state**?
- Challenge: bottleneck link is remotely located

Steady state: Ideal goal

- **High sending rate:** Use the full capacity of the bottleneck link
- **Low delay:** Minimize the overall delay of packets to get to the receiver
 - Overall delay = propagation + queueing + transmission
 - Assume propagation and transmission components fixed
- “Low delay” reduces to **low queueing delay**
- i.e., don’t push so much data into the network that packets have to wait in queues
- Key question: When to send the next packet?

When to send the next packet?



Rationale


- When the sender receives an ACK, that's a signal that the previous packet has left the bottleneck link (and the rest of the network)
- Hence, it must be safe to send another packet without congesting the bottleneck link
- Such transmissions are said to follow packet conservation
- ACK clocking: “Clock” of ACKs governs packet transmissions

ACK clocking: analogy


- How to avoid crowding a grocery store?
- Strategy: Send the next waiting customer exactly when a customer exits the store
- However, this strategy alone can lead to inefficient use of resources...



ACK clocking alone can be inefficient



ACK clocking alone can be inefficient



The sending rate should be high enough to “fill up” the link
Analogy: a grocery store with only 1 customer in entire store
If the store isn't “full”, you're using store space inefficiently

Steady State of Congestion Control

- Send at the highest rate possible (to fully use the link)
- while being ACK-clocked (to avoid congesting the pipe)
- So, how to get to steady state?

Finding the Right Congestion Window

Let's play a game


- Suppose I'm thinking of a number (positive integer). You need to guess the number I have in mind.
- Each time you guess, I will tell you whether your number is smaller or larger than (or the same as) the one I'm thinking of
- My number can be very large or small
- How would you go about guessing the number?

Finding the right congestion window

- TCP congestion control algorithms solve a similar problem!
- There is an **unknown** bottleneck link rate that the sender must match
- If sender sends more than the bottleneck link rate:
 - packet loss, delays, etc.
- If sender sends less than the bottleneck link rate:
 - all packets get through; successful ACKs


Quickly finding a rate: TCP slow start

- Initially $cwnd = 1 \text{ MSS}$
 - MSS is “maximum segment size”
- Upon receiving an ACK of each MSS, increase the $cwnd$ by 1 MSS
- Effectively, double $cwnd$ every RTT
- Initial rate is slow but ramps up **exponentially fast**
- On loss (RTO), restart from $cwnd := 1 \text{ MSS}$



Behavior of slow start

Assume bottleneck link (rate) can change any time. Hence, keep repeating the process from the start.



Slow start has problems

- Congestion window **increases too rapidly**
 - Example: suppose the “right” window size `cwnd` is 17
 - `cwnd` would go from 16 to 32 and then dropping down to 1
 - Result: massive packet drops
- Congestion window **decreases too rapidly**
 - Suppose the right `cwnd` is 31, and there is a loss when `cwnd` is 32
 - Slow start will resume all the way back from `cwnd` 1
 - Result: unnecessarily low throughput
- Instead, perform **finer adjustments** of `cwnd` based on signals

Use slow start mainly at the beginning

- You might accelerate your car a lot when you start, but you want to make only small adjustments after.
 - Want a smooth ride, not a jerky one
- Slow start is a good algorithm to get close to the bottleneck link rate when there is little info available about the bottleneck, e.g., the beginning of a connection
- Once close enough to the bottleneck link rate, use a different set of strategies to perform smaller adjustments to the congestion window
 - Called TCP congestion avoidance