Transmission Control Protocol (TCP)

- Reliable
- Connection-oriented
- Point-to-point
- Full-duplex
- Streams, not messages

Initialization: 3 Way Handshake

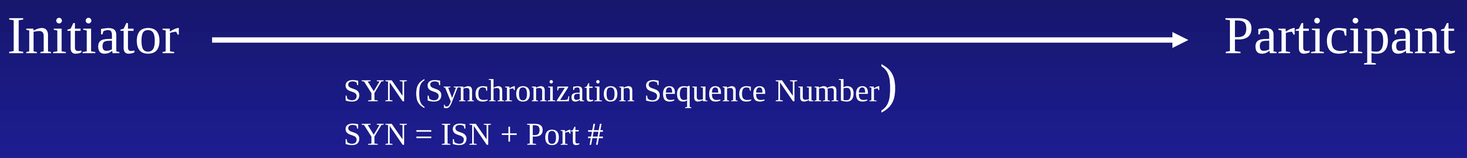

Initiator → Participant

SYN (Synchronization Sequence Number)
SYN = ISN + Port #

- The client begins it's active open by sending a SYN to the server. SYN stands for "Synchronization Sequence Number", but it actually contains much more.

- The SYN message contains the initial sequence number (ISN). This ISN is the starting value for the sequence numbering that will be used by the client to detect duplicate segments, to request the retransmission of segments, &c.

- The message also contains the *port number*. Whereas the hostname and IP address name the machine, the port number names a particular processes. A process on the server is associated with a particular port using bind().

Initialization: 3 Way Handshake

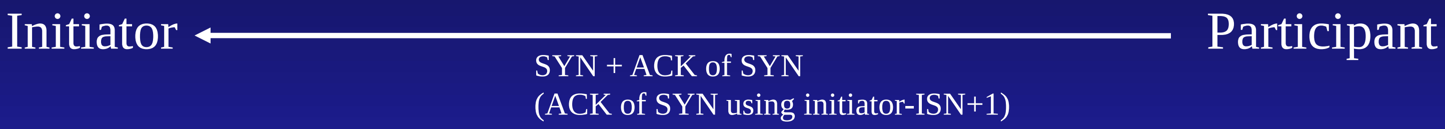

Initiator

Participant

SYN + ACK of SYN
(ACK of SYN using initiator-ISN+1)

- The server performs the passive open, by sending its own ISN to the client. It also sends an Acknowledgement (ACK) of the client's SYN, using the ISN that the client sent plus one.

Initialization: 3 Way Handshake

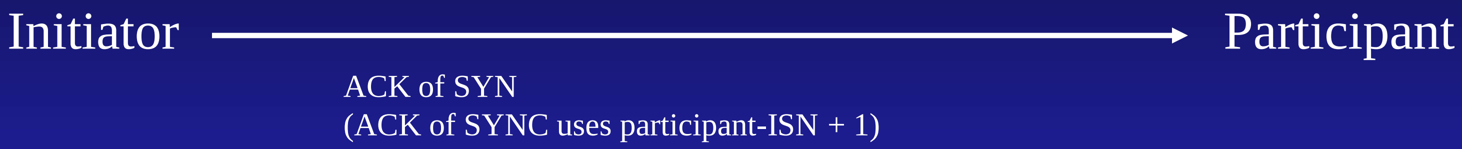

Initiator → Participant

ACK of SYN
(ACK of SYNC uses participant-ISN + 1)

- The last step is for the client to acknowledge the server’s SYN

Initialization: 3 way Handshake

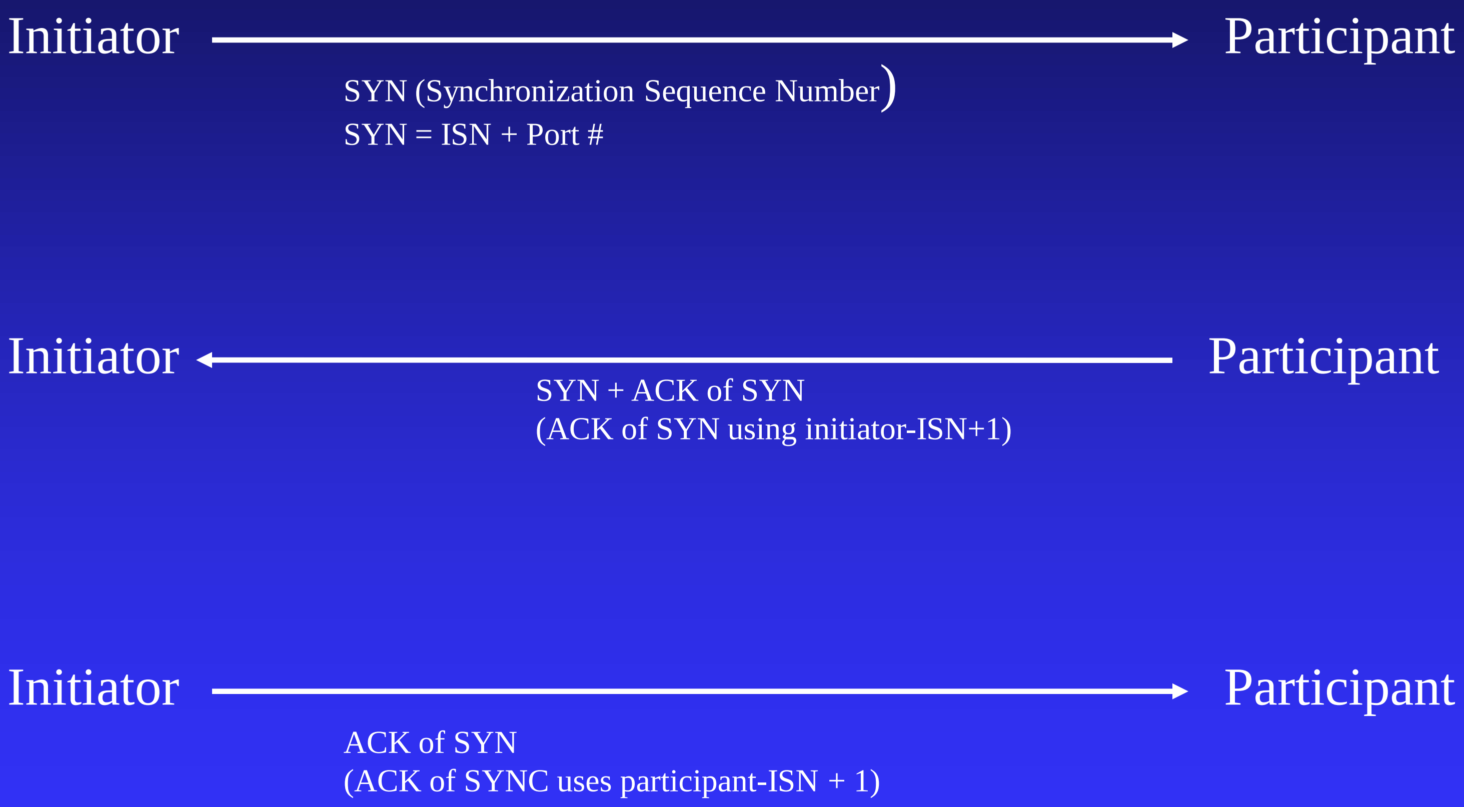

Initiator → Participant

SYN (Synchronization Sequence Number)

SYN = ISN + Port #

Initiator ← Participant

SYN + ACK of SYN

(ACK of SYN using initiator-ISN+1)

Initiator → Participant

ACK of SYN

(ACK of SYNC uses participant-ISN + 1)

How and Why is the ISN Chosen?

- Why do we send the ISN, instead of just always start with 1?
- The answer to this is that we don't want to misinterpret an old segment. For example, consider a short-lived client process that always talked to the same server. If the ISN's would always start with one, a delayed segment from one connection might be misinterpreted as the next segment for a newer instance of the same client/server-port combination. By doing something more random, we reduce the bias toward low sequence numbers, and reduce the likelihood of this type of situation.
- RFC 793 specifies that the ISN should be selected using a system-wide 32-bit counter that is incremented every 4 microseconds. This approach provides a "moving target" that makes segment number confusion unlikely.
- 4.4BSD actually does something different. It increments the counter by 64K every half-second and every time a connection is established. This amortizes to incrementing the counter by one every 8 microseconds.

Connection Termination

- When either side of a TCP connection is done sending data, it sends a FIN (finished) to the other side. When the other side receives the FIN, it passes an EOF up the protocol stack to the application.
- Although TCP is a full-duplex protocol, the sending of a FIN doesn't tear down the whole connection. Instead it simply indicates that the side sending the FIN won't send any more data. It does not prevent the other side from sending data. For this reason, it is known as a *half-close*. In some sense, a half-closed connection is a half-duplex connection.
- Although TCP allows for this half-closed state, in practice, it is very rarely used. For the most part, when one side closes a connection, the other side will immediately do the same. It is also the case that both sides can concurrently sends FINs. This situation, called a *simultaneous close* is perfectly legal and acceptable.

One Side

Other side

ACK of SYN
(ACK of SYNC uses participant-ISN + 1)

Half Close

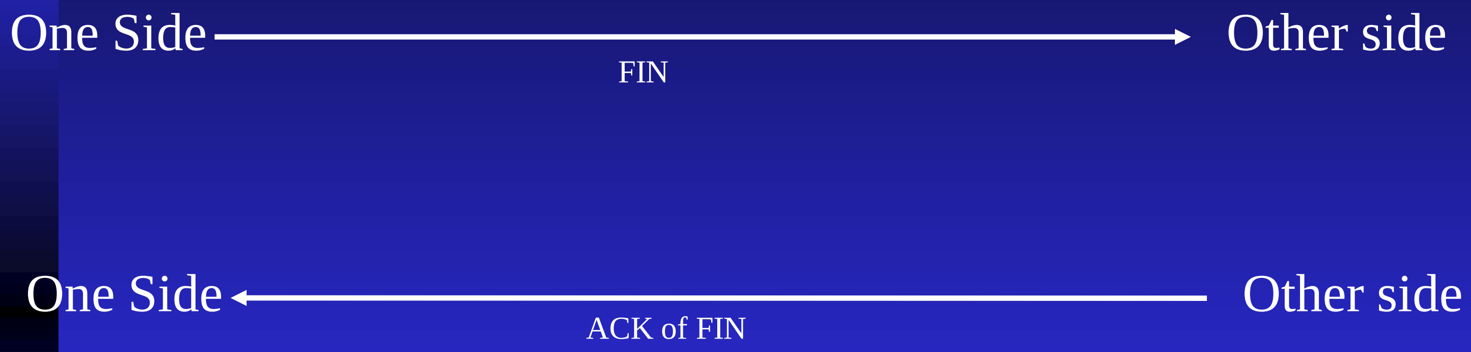

One Side

FIN

Other side

One Side

ACK of FIN

Other side

Maximum Segment Life

- MSL stands for *Maximum Segment Life*.
- Basically, MSL is a constant that defines the maximum amount of time that we believe a segment can remain in transit on the network.
- 2MSL, twice this amount of time, is therefore an approximation of the maximum round trip time.
- We wait 2MSL after sending the ACK of the FIN, before actually closing the connection, to protect against a lost ACK.
- If the ACK is lost, the FIN will be retransmitted and received. The ACK can then be resent and the 2MSL timer restarted.

What About Crashes, &c.

- But wait, if both sides need to close the connection, what happens if the power fails on one side? Or a machine is shut off? Or the network goes down?
- Well, the answer to this is very simple: Nothing. Each side will maintain at least a half-open connection until the other side sends a FIN. If the other side never sends a FIN, barring a reboot, the connection will remain at least half-open on the other side.
- What happens if neither process ever sends data? The answer to this is also very simple: Nothing. Absolutely nothing is sent via TCP, unless data is being sent.

TCP Keep-Alive Option

- Well, some people were as upset as you were by the idea that a half-open connection could remain and consume resources forever, if the other side abruptly died or retired. They successfully lobbied for the *TCP Keepalive Option*.

- This option is disabled by default, but can be enabled by either side. If it is enabled on a host, the host will probe the other side, if the TCP connection has been idle for more than a threshold amount of time.

- This timer is system-wide, not connection wide and the RFC states that, if enabled, it must be no less than two hours.

- Many people (including your instructor) believe that this type of feature is not rightfully in the jurisdiction of a transport layer protocol. We argue that this type of session management is the rightful jurisdiction of the application or a session-level protocol.

- Please do realize that this is a religious issue for many and has received far more discussion than it is probably worth. Independent of your beliefs, please don't forget that the timer is system-wide -- this can be a pain and might even lead many keepalive-worshipers opt for handling this within the applications.

Reset (RST)

- TCP views connections in terms of *sockets*. A popular author, Richard Stevens refers to these as *connections* -- this is wrong, but has worked its way into the popular vernacular.
- A socket is defined as the following tuple:

  <destination IP address, destination port #, source IP address, source port number>

- A RST is basically a suggestion to abort the connection.
- A reset will generally be sent by a host if it receives a segment that doesn't make sense. Perhaps the host crashed and then received a segment for a port that is no longer in use.
- In this case, the RST would basically indicate, "No one here, but us chickens" and the side that received the RST would assume a crash, close its end and roll-over or handle the error.

Transferring Data

- TCP operates by breaking data up into pieces known as *segments*.
- The TCP packet header contains many pieces of information. Among them is the Maximum Segment Length (MSL) that the host is willing to accept.
- In order to send data, TCP breaks it up into segments that are not longer than the MSL.

Acknowledgement

- Fundamentally, TCP sends a segment of data, including the segment number and waits for an ACK. But TCP tries to avoid the overhead involved in acking every single segment using two techniques.
- TCP will wait up to 200mS before sending an ACK. The hope is that within that 200 mS a segment will need to be sent the other way. If this happens, the ACK will be sent with this segment of data. This type of ACK is known as a *piggyback ACK*.
- Alternatively, no outgoing segment will be dispatched for the sender within the 200mS window. In this case the ACK is send anyway. This is known as a *delayed ACK*.
- *Note*: My memory is that the RFC actually says 500mS, but the implementations that I remember use a 200mS timer. No big deal, either way.

More About the ACKs

- By default, TCP uses *cumulative acknowledgement* unless a segment arrives out of order, in which case TCP will use an *immediate acknowledgement* of the last contiguous segment received. This tells the sender which segment is expected. This is based on the assumption that the likely case is that the missing segment was lost not delayed.

  If this assumption is wrong, the first copy to arrive will be ACKed, the subsequent copy will be discarded.

- *Selective acknowledgement* is an option that can be negotiated by the sender and receiver if both support it. Instead of sending an integer representing a single number, it uses the bits at a bit map, where each bit represents a segment within the window: *1* for received, *0* for absent. This gives the sender a much better understanding of which segments need to be resent, mitigating the need to sent resent segments already received, but located after the last contiguous segment.

Nagle Algorithm

- One interesting observation is that it takes just as much overhead to send a small amount of data, such as one character, as it does a large amount of data, such as a full MSL of data.
- The massive overhead associated with small segments can be especially wasteful if the network is already bogged down.
- One approach to this situation is to delay small segments, collecting them into a full segment, before sending. This approach reduces the amount of non-data overhead, but it can unnecessarily delay small segments if the network isn't bogged down.
- The compromise approach that is used with TCP was proposed by Nagle. The Nagle Algorithm will send one small segment, but will delay the others, collecting them into a larger segment, until the segment that was sent is acknowledged. In other words, the Nagle algorithm allows only one unacknowledged small segment to be send.

Nagle Algorithm

- This approach has the following nice property. If the network is very bogged down, the ACK will take a long time. This will result in many small segments being collected into a large segment, reducing the overhead. If the network isn't bogged down, the ACK will arrive very rapidly, allowing the next small segment to be sent without much delay. If the network is fast, fewer small segments will be concatenated, but who cares? The network isn't doing much else.

- In other words, the Nagle algorithm favors the sending of short segments on a "fast network" and favors collecting them into larger segments on a "slow network." This is a very nice property!

- There are certain circumstances where the Nagle approach should be disabled. The classic example is the sending of mouse movements for the X Window system. In this example, it is critically important to dispatch the short packets representing mouse movements in a timely way, independent of the load on the network. These packets need a response in soft real-time to satisfy the human user.

The Sliding Window Model

- As we mentioned earlier, TCP is a sliding window protocol much like the example protocol that we discussed last class. The sliding window model used by TCP is almost identical to model used in the example.
- In the case of TCP, the receiver's window is known as the *advertised window* or the *offered window*. The side of the window is advertised by the receiver as part of the TCP header attached to each segment. By default, this size is usually 4096 bytes.
- The *usable window* is the portion of the advertised window that is available to receive segments.
- The only significant difference is the one that we mentioned before: TCP uses a cumulative ACK instead of a bit-mask.

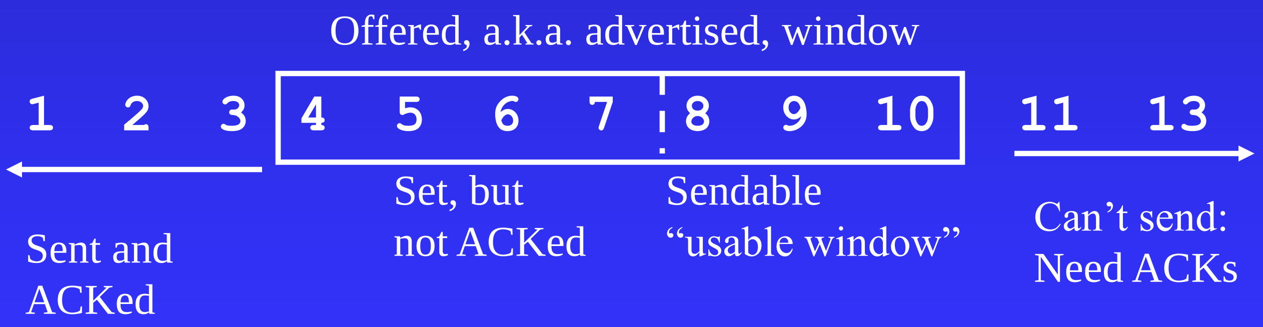

Offered, a.k.a. advertised, window

1 2 3 4 5 6 7 8 9 10 11 13

Sent and ACKed

Set, but not ACKed

Sendable “usable window”

Can’t send: Need ACKs

Slow Start and Congestion Avoidance

- The advertised window size is a limit imposed by the receiver. But the sender doesn't necessarily need or want to send segments as rapidly as it can in an attempt to fill the receiver's window.

- This is because the network may not be able to handle the segments as rapidly as the sender can send them. Intermediate routers may be bogged down or slow. If the sender dispatches segments too rapidly, the intermediate routers may drop them requiring that they be resent.

- In the end, it would be faster and more bandwidth efficient to send them more slowly in the first place.

- TCP employs two different techniques to determine how many segments can be sent before acknowledgement: *slow start* and *congestion avoidance*.

- These techniques make use of a sender window, known as the *congestion window*. The congestion window can be no larger than the receiver's advertised window, but may be smaller. The congestion window size is known as *cwnd*.

Slow Start

- Initially, the congestion window is one segment large. The sender will send exactly one segment and wait for an acknowledgement.

- Then the sender will send two segments. Each time an ACK is received, the congestion window will grow by two. (This results in 1,2,4,8,16,… growth)

- This growth will continue until the congestion window size reaches the smaller of a threshhold value, *ssthresh* and the advertised window size.

- If the congestion window reaches the same size as the advertised window, it cannot grow anymore.

- If the congestion window size reaches ssthresh, we want to grow more slowly -- we are less concerned about reaching a reasonable transmission rate than we are about suffering from congestion. For this reason, we switch to congestion avoidance.

- The same is true if we are forced to retransmit a segment -- we take this as a bad sign and switch to congestion avoidance.

Congestion Avoidance

- Congestion avoidance is used to grow the congestion window slowly.
- This is done after a segment has been lost or after ssthresh has been reached.
- Let's assume for a moment that ssthresh has been reached. At this point, we grow the congestion window by the greater of 1 segment and (1/cwnd). This rate or growth is slower than it was before, and is more appropriate for tip-toeing our way to the network's capacity.

```
cwnd = cwnd + MAX (1, (1/cwnd))
```

Congestion Avoidance

- Eventually, a packet will be lost. Although this could just be bad luck, we assume that it is the result of congestion -- we are injecting more packets into the network than we should.
- As a result, we want to slow down the rate at whcih we inject packets into the network. We want to back off a lot, and then work our way to a faster rate. So we reset ssthresh and cwnd:

```
ssthresh = MAX (2, cwnd/2)
cwnd = 1
```

After Congestion Avoidance

- After reducing the congestion window, we reinvoke slow start.
- This time it will start with a cwnd size of 1 and grow rapidly to half of the prior congestion window size. At that point congestion avoidance will be reinvoked to make tip-toe progress toward a more rapid transmission rate.
- Eventually, a packet will be lost, ssthresh will be cut, cwnd will be reset to 1, and slow start will be reinvoked.
- It is important to notice that ssthresh doesn't always fall -- it can grow. Since ssthresh is set to (cwnd/2), if the new value of cwnd is more than twice the old value of ssthresh, ssthresh will actually increase.
- This makes sense, because it allows the transmission rate to slow down in response to a transient, but to make a substantial recovery rapidly. In this respect, the exponential growth rate of "slow start" is actually a "fast start".

An Example of Slow Start and Congestion Avoidance

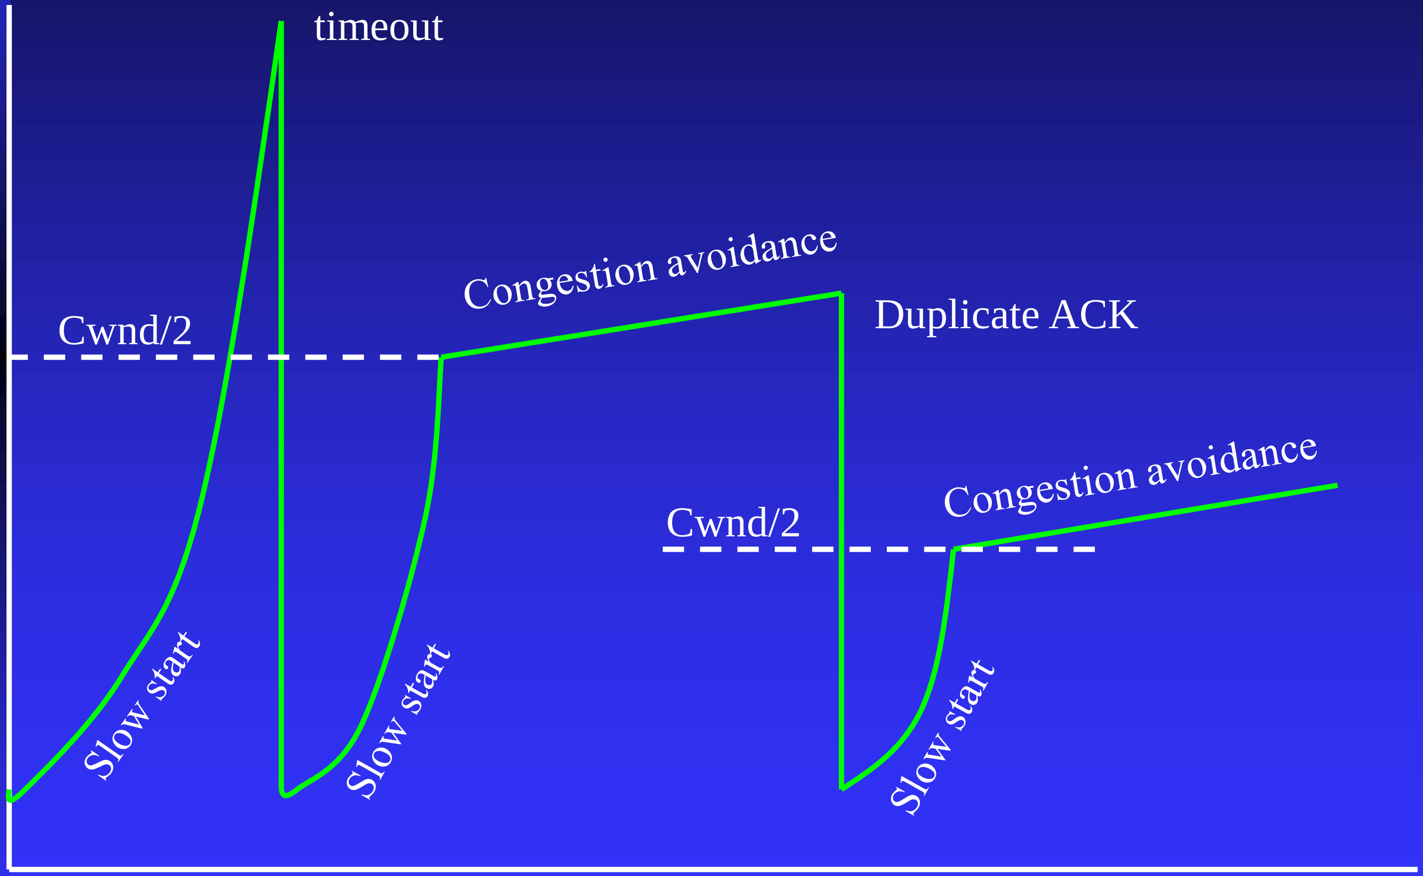

More Tweaking

- It is clear that the traditional TCP slow start and congestion avoidance schemes assume that a single missing AC indicates congestion. But, it could actually represent just dumb, bad luck – and impose a really nasty penalty upon being unlucky.

- There are many newer variants, also options, that impose lesser penalties earlier, in an attempt to balance the possibility of bad luck against the harm caused by pumping segments into a suffocating network.