Traceroute and TCP Tahoe: Congestion Control

Joshua Tan

For my capstone course, I took CSCI1680 (Computer Networks). On top of the normal coursework, I had to implement additional features for two of the major projects, TCP and IP. For TCP, I implemented an additional congestion control algorithm and for IP, I implemented traceroute.

1 Congestion Control - TCP Tahoe

Congestion control is a concept where one uses an algorithm to send packets at the most optimal rate. In a network, there could be choke points where the incoming traffic exceeds the outgoing bandwidth. This leads to packets being lost when too much traffic is sent through these choke points. Thus, we do not want to send packets at a rate that is too high as packet loss will lead to re-transmissions, taking even more time under the TCP protocol. Additionally, we also don't want to set a rate that is too low because we would be underutilizing the network and we could be improving the network speed. Congestion control, along with other improvements, allows us to find the most optimal packet rate that leads to the fastest transmission of data across a network.
For my congestion control algorithm, I chose TCP Tahoe. It is a congestion algorithm that utilizes slow start, congestion avoidance, and fast re-transmit. Without getting too deep into the technical details, slow start and congestion avoidance are the mechanism in which the algorithm "discovers" the most optimal packet rate. Fast re-transmit is a mechanism that allows for faster re-transmission of dropped packets compared to the vanilla TCP protocol. By retransmitting based on duplicate acknowledgements instead of a timeout, lost packets are detected much quicker, allowing for quicker retransmissions.

2 IP Traceroute

For the IP project, I implemented traceroute. Traceroute is a network diagnostic tool used to track the path taken by a packet on an IP network from source to destination, reporting the IP addresses of the nodes it went through. This is done by sending packets of various TTLs (time to live). When a node receives a packet that has a TTL of 0, it sends an ICMP packet to the original source. By sending out packets with various TTLs, we can determine the route a packet takes based on the ICMP response packets received.