Position paper on CC for Interactive RT Communication

Matt Mathis

It is important that any congestion control algorithm or protocol chosen to support Real Time communication include instrumentation suitable for making the case that the network needs to support some mechanism to protect delay sensitive traffic from throughput maximizing traffic. Furthermore all applications should include an instrumentation and diagnostic console to display these instruments.

It is well understood that in some environments (e.g. when the entire path has idle capacity) the RT congestion control problem can be solved fairly easily. It is also understood that in other environments the problem is hopeless in the sense that the cross traffic or some other network phenomena will cause unacceptable delay or delay jitter, no matter what the RTCC algorithm does. A very sophisticated RTCC algorithms might be able to slightly move the partition between these two cases, but even heavy investment in in such technologies will at best only marginally increase the applicability RT protocols.

The problem is that TCP requires a queue that is least some fraction of the RTT in order to fill the link (details determined by the specific implementation and version). If RT traffic is sharing a queue with TCP, the RT traffic will be delayed by the queues caused by TCP. These queues are controlled by the drop properties of the link or AQM. If the RTT is sufficiently large the minimum queue required for TCP to fill the link will be larger than the largest queue that the RT traffic can tolerate.

We conclude from this situation that even omniscient AQM would be insufficient. The only solution is to segregate RT traffic from TCP (and other throughput maximizing traffic) such that they don't share queues. This implies that the Internet needs end-to-end QoS or some other mechanism to implement this segregation.

Given that end-to-end QoS has proven to be hard to deploy, it is imperative that RT applications have built into them mechanisms to give stakeholders insights into why they don't work as well as desired, inorder to motivate the community to invest in deploying QoS or alternatives.

My claim is that a modest investment in instrumentation will have more long term impact than a huge investment in complicated algorithms to squeeze the last bit of capacity out of what is fundamentally an unsolvable problem.

mattmathis@google.com