

## ABSTRACT

### **MARTIN, JAMES JOHN. RTT-Based Congestion Avoidance for High Speed TCP Internet Connections. (Under the direction of Arne Nilsson and Injong Rhee.)**

The Internet provides several forms of congestion avoidance. TCP's slow-start and congestion avoidance algorithm have been crucial to ensuring the stability of the Internet. Additional end-to-end congestion avoidance algorithms that use increases in round trip time (RTT) samples as an implicit congestion feedback indication have been proposed (e.g., TCP/Vegas, TCP/Dual). The motivation and claimed benefits of such algorithms are that they can offer performance improvements at both the endpoint (i.e., higher throughput) and within the network (i.e., less buffers consumed during times of congestion). Furthermore, as the algorithms are confined to the sender, the improvement can be incrementally deployed in today's best effort Internet.

Our research focuses on end-to-end congestion avoidance algorithms that use round trip time (RTT) fluctuations as an indicator of the level of congestion over the path. The algorithms are referred to as *delay-based congestion avoidance* or DCA. In particular, we are interested in an incremental enhancement to the TCP/Reno protocol that uses additional implicit congestion feedback derived from a per packet RTT measurement. It can take years for Internet protocols to standardize and even longer for Internet standards to be widely deployed. Consequently, incrementally deployable enhancements, especially to TCP, are of high interest to the Internet community. An incrementally deployable enhancement to TCP should meet the following requirements:

- It must improve the throughput of the TCP connection that employs the enhancement.
- It should not reduce the performance of other competing TCP flows on the same path where the "enhanced" TCP flow travels.
- Ideally it requires changes only to a TCP sender.

Over a best effort Internet connection where TCP/Reno flows dominate, RTT-based congestion avoidance can only improve TCP throughput if the algorithm is able to reduce the frequency of packet loss. We

evaluate the effectiveness of DCA in predicting and avoiding packet loss over high speed Internet paths to improve TCP throughput. Previous studies of DCA algorithms (e.g., TCP/Vegas and TCP/Dual) have focused on lower speed paths. Because technology continues to increase link speeds, we are interested in high speed paths where the lowest link capacity along the path is 10 mbps. We also assume that a DCA flow will consume only a fraction of the total bandwidth. Recent measurement studies of Internet traffic confirms that Internet backbone switches, which generally support 45mbps links (or faster), are subject to thousands of low bandwidth TCP flows. Furthermore, it has been observed that packet loss rates over certain portions of the Internet are increasing. The objective of DCA in this environment is to avoid packet loss in an attempt to increase the utilization of available bandwidth.

The thesis of this dissertation is to show that RTT-based congestion avoidance cannot be reliably incrementally deployed over high speed Internet paths. We prove our thesis through measurement analysis using actual TCP traces collected over five days, and simulation studies of DCA algorithms using the Network Simulator (*ns*) whose traffic dynamics are modeled after the actual Internet paths. Our analysis shows two findings:

- A TCP sender that is extended with DCA will typically result in degraded throughput primarily because the congestion information contained in RTT samples cannot be reliably used to predict packet loss.
- By design, a DCA algorithm will limit the number of packets allowed in the network during periods of congestion (as compared to a TCP/Reno connection). When the level of traffic generated by a single TCP connection over a high speed path represents a small fraction of the total traffic that flows over the path, the reaction of a single DCA flow will have minimal impact on the congestion level over the path.