6 Summary of Measurement Analysis

In the preceding chapters, we have presented evidence that the relationship between increases in the TCP packet RTT samples and loss events that exist over high speed Internet paths are not conducive to DCA. Specifically our measurement results show:

- The correlation indication metric shows that loss is preceded by an increase in RTT on the order of 30-50% of the time. Of the loss events that are preceded by an increase in RTT, only 7-18% reflect a significant increase in RTT (i.e., where the increase is greater than a long term average of past RTT samples plus a threshold of one standard deviation). This will lead to a low number of DCA congestion decisions that have a chance of actually avoiding loss events.
- The loss conditioned delay correlation metric also shows that there is a weak level of correlation between an increase in RTT samples and loss. The level of correlation appears to be stronger over certain paths. The metric also suggests that packets that are sent closer than one RTT to the time that the dropped packet is sent have a higher probability of reflecting an increase in RTT (i.e., 50-70% of the time). This indicates that the queue buildup associated with a significant percentage of the loss events that we observed (perhaps up to 20%) is of time scale less than one RTT. The implication is that by the time DCA detects an increase in RTT (because the feedback loop time is large) any reaction would be too late to avoid the loss.
- The *tcpRTT* CDF shows that increases in RTT are common. The loss conditioned delay CDF metric shows that loss generally occurs above some minimum RTT threshold level. Above the threshold, the value of the RTT prior to loss will generally follow the same distribution as the *tcpRTT* samples. In other words, the increase in RTT that accompanies loss is indistinguishable from the other more frequent increases in RTT. The implies that by attempting to avoid the 7-18% of loss events that are potentially predictable, many additional congestion reactions will also occur, which may detrimentally affect the throughput of DCA.

These results indicate that the congestion information contained in TCP packet RTT samples may be insufficient to allow an endpoint to reliably predict future packet loss events. We conjecture that two reasons for this are:

- A TCP constrained RTT sampling process is insufficient to track the bursty congestion associated with loss over high speed paths.
- An end-to-end congestion sampling process is not able to accurately assess the congestion level of a bottleneck router.

To see the impact of DCA reactions to the overall throughput of a DCA connection, we feed an analytic TCP/DCA throughput model the data based on measurements. Our results show that the DCA algorithm will degrade TCP throughput in the range of 10% to 60% as compared to a similar TCP/Reno connection. This is true even if the DCA algorithm is able to avoid a high percentage of the actual packet loss events. These results suggest that the predicted throughput degradation is much greater over the lower loss rate paths, especially those exhibiting significant RTT variations. Intuitively, this makes sense as the DCA algorithm will react frequently with little chance of actually avoiding a loss event. For higher loss situations, the congestion control is dominated by TCP's loss recovery algorithms reducing the impact of the DCA control.

The throughput model analysis relies on the following assumptions:

- The throughput model analysis based on the TCP throughput model assumes that a congestion reaction by a DCA flow will not alter the congestion dynamics that are active over the path
- We assume that a hypothetical DCA algorithm has the following properties:
  - Reacts to an increase in RTT by reducing the TCP *cwnd* by half.
  - Reacts once per RTT or once per congestion epoch.
  - Limited to specific moving window sizes.

We consider these assumptions in more detail in the next chapter of the dissertation.