10 Summary of Simulation Analysis

In the preceding chapters, we have presented a simulation analysis that allowed us to confirm that an RTT does not provide enough information for an endpoint to reliably predict packet loss events. The simulation analysis extends the measurement analysis which was limited to an analysis based on end-to-end probes. Only by observing the internal queue dynamics were we able to show that the following are the fundamental problems.

- A TCP constrained RTT sampling process is not able to track the bursty congestion associated with loss over high speed paths. The limiting factors are:
  - ♦ The rapid queue growth prior to packet loss with respect to the RTT. The simulation run confirms that the rate of queue increase at a congested router can be less than a RTT making it impossible for a DCA algorithm to react in time to prevent loss.
  - ♦ For cases when the rate of queue increase prior to loss is one the order of 1-2 RTT's, the coarseness of a TCP constrained RTT probe impacts the ability of a DCA to reliably detect the increase in RTT in time to prevent the loss. The coarseness of the sampling process is affected by the following:
    - ▪ The number of probe packets is limited by TCP's congestion control algorithms. So, during times of congestion, the number of probes tends to reduce.
    - ▪ Because of TCP's bursty behavior, probes tend to be sent "clumped" together every RTT.
    - ▪ As the queue size increases, the congestion indication feedback time grows with the increase in RTT further affecting the sampling process.
- Aside from the sampling difficulties described above, other factors contribute to making increases in RTT an unreliable indicator of future loss events. The measurement results indicates that there are many RTT variations and that loss is not significantly correlated with large RTT values. We have shown that one explanation for this is when loss occurs at multiple links with each provisioned differently (i.e., different link speed and buffer capacities). For example, loss at a 155mbps link configured with a maximum buffer size of 200 packets leads to a queue delay that is insignificant

compared with the delays associated with a 45 mbps router configured with a maximum buffer size of 600 packets.

Using the TCP/DCA simulation model, we confirm our assumption that the reactions by a DCA flow do not significantly reduce the congestion level of the bottleneck links. It was necessary to show this to support our claim that there is no benefit to incrementally deploying TCP/DCA. To validate the throughput analysis from Part 1 of the dissertation, we showed that the congestion dynamics at a bottleneck link are not significantly impacted by the different behavior caused when a TCP/Reno flow is replaced with a TCP/DCA flow over the path. We found that when the competing traffic at a congested link consists primarily of low bandwidth TCP flows, the queue oscillations are not significantly impacted by the DCA congestion reactions. More importantly, we found that the dynamics associated with the congestion (e.g., loss rates and average queue delays) are not impacted.

Our objective was to show that DCA cannot be incrementally deployed over high speed Internet paths. The measurement analysis was restricted to a particular algorithm. To validate that the results hold over all DCA algorithms, we performed the following analysis:

- We ran simulation experiments showing that DCA will not improve throughput for a range of algorithm parameters that included:
    - ♦ The amount of send rate reduction in response to an increase in RTT.
    - ♦ The maximum allowed frequency of congestion reaction.
    - ♦ The window size associated with the moving window averages used by DCA.
- We ran similar simulation experiments using models of the TCP/Vegas and TCP/Dual algorithms and find that they also will not improve the throughput of the connection. We observe that certain aspects of each of these algorithms hinders the abilities of the DCA algorithm to avoid packet loss. For example, the congestion detection process used by Dual, in particular the level of control applied by the algorithm, is highly dependent on the dynamics of the congestion (i.e., it depends on the maximum RTT sample). We have seen that the enhanced loss recovery algorithm of Vegas can be effective at improving TCP throughput. By isolating the CAM algorithm we have observed that,

because of its linear congestion reaction and its RTT filtering properties, CAM is more likely to react to long term congestion than to brief RTT fluctuations. And we have seen that the level of control exerted by Vegas will actually decrease as the congestion level increases.