Chapter 7

Conclusion

The analysis, simulations, and experiments carried out in this work substantiates and elaborates on the suggested *tiny buffers* rule for Internet backbone routers [25]: backbone routers perform well with buffers of size 10-50 packets if their arrival traffic is non-bursty. We saw (in a wide variety of simulation and experimental settings) that if a bottleneck link runs a few tens of times faster than the access links, then the traffic will be sufficiently smooth to achieve over 80% link utilization with tiny buffers. In our simulations and experiments we assumed a core-to-access bandwidth ratio of 10-100. Our measurements on commercial backbone links, however, show that the vast majority of flows come from much slower access links; this could result in even higher utilization.

From the users' perspective, short flows could benefit from tiny buffers and experience faster completion times. This is because of a shorter round-trip time in a network with tiny buffers: reducing the buffer size to a few packets almost eliminates queueing delays. Larger flows could experience a longer completion time with tiny buffers if the link is heavily congested. A sufficiently over-provisioned bottleneck link eliminates this increased download time (Chapter 4).

The tiny buffers result is not limited to a single router. We can implement such buffers in all routers inside a backbone network as long as the ingress traffic is smooth. In general, we don't expect to see a difference between traffic burstiness at ingress ports and inside the backbone network. For cases where adversarial traffic matrix or network topology make the ingress traffic bursty, we proposed Bounded Jitter Policy (Chapter 5), which makes the traffic at each router behave as if it comes directly from the ingress ports of the network.

While our simulations and experiments strongly suggest that routers perform well with tiny buffers, ultimately, network operators and service providers need to verify if it is so in operational backbone networks. When verified, this can have a significant implication in building all-optical routers, which have very limited buffering capacities.

Most of the simulations and experiments run in this work are with TCP Reno implemented at the end hosts. As we saw in Chapters 3 and 4, more recent variants of TCP, which are designed for large delay-bandwidth networks, achieve even higher throughput with tiny buffers compared to TCP Reno. Buffer sizing for traffic that uses other (non-TCP) congestion-aware algorithms requires further study.

The crucial assumption in the results of this work is that the arrival traffic is non-bursty. Buffer sizing in networks with different traffic patterns (e.g., data centers) needs to be studied separately. The problem has attracted interests in data center design because low-cost commodity switches, suggested as the connecting fabric in these networks [8], have small buffering capacities.