Chapter 6

Conclusions and Future Work

In this thesis, we have presented techniques for maximizing the utilization of a shared resource, network bandwidth. For the studies we performed, we focused on the largest consumers of this resource, bulk transfer applications, in which a large volume of data has to be reliably transmitted from a single source to one or more destinations. The first technique we considered established a rate-based allocation of end-to-end bandwidth to bulk connections so as to maximize the aggregate throughput of the connections, measured as the sum of the transmission rates. Starting from related work which proved tight bounds on how well a set of point-to-point connections could approximate this global objective using only local information, we explored the tradeoff between the amount of communication and the quality of the approximation. Our main result shows that a small (polylogarithmic) number of distributed communication rounds suffice to generate a near-optimal $(1+\epsilon)$ approximation to the optimum solution. Moreover, extensions of this distributed technique apply to allocations in which reliable multicast and layered multicast connections participate. Rate-based allocation techniques with provable performance guarantees provide an alternative to standard congestion control techniques, in which connections make dynamic, and often rapid adjustments to their transmission rates based on the presence or absence of perceived congestion along their route.

The second technique we addressed focuses on the on-line scheduling problem of determining what data to place in each packet of a bulk transfer to maximize the effective throughput at the receivers throughout the transmission. The policies we develop for this problem can be used in conjunction with the rate-based allocation policies we develop, or in conjunction with any other mechanism for regulating the rate at which connections inject packets into the network. Factors shaping the development of the Internet motivate the design of policies which deliver high performance even as the impact of other users in the network causes transmission rates to fluctuate, round-trip times to vary widely and packets to be lost. For the unicast case, one of our main contributions is a randomized data selection policy which delivers provably high performance with respect to competitive measures on arbitrary event sequences. This result further implies that data selection policies need not be tuned for ideal performance with a given rate allocation policy – the algorithm we present will deliver near-optimum performance when coupled with *any* rate allocation policy. This provides further motivation for the modular separation we advocate between policies which regulate the transmission *rate* and policies which regulate the *content* of transmissions.

For multicast applications, we show that the use of coding techniques, and forward error correction in particular, is a prerequisite for a multicast data selection policy to deliver provably high performance. The use of forward error correction has an additional benefit in that it can dramatically reduce the number of transmissions needed to complete a reliable bulk transfer to multiple receivers when packet loss is present. Moreover, the gains which are realized over strategies which use packet retransmission grow with the number of receivers in the multicast group.

6.1 Future Work

For the problem of rate allocation, there are several avenues along which to perform further exploration. One consideration is the discrepancy between the performance predicted by the analysis and the performance realized by the actual implementation. For the examples we have tested, the algorithm realizes a considerably closer approximation ratio than the target ratio. For a given setting of ϵ, realization of an $(1+\frac{\epsilon}{10})$ approximation is quite common. This has clear implications for the specification of ϵ in practice, and it would be nice to tighten the analysis or exhibit examples on which the current analysis is tighter. A related problem regards the distributed running time – for all examples on which we have run the algorithm, the dual feasible solutions achieve the desired approximation ratio after only 1/10th of the run is complete. A natural question would be to ask if there is a way to exploit this behavior and improve the running time, either by taking less conservative steps while increasing primal variables, or by improving the mechanism which correlates values of primal and dual variables.

Another direction would be to consider generalizing the objective functions over which we perform the optimization. The weighted throughput version of the problem can be viewed as a profit maximization problem on pay-per-use networks in which connections are willing to pay a fixed amount for each unit of end-to-end rate which they are allocated. But in many economic settings, of which this is an example, the consumer is likely to derive diminishing marginal benefit for each additional unit of commodity beyond some threshold. This motivates study of a related non-linear optimization problem in which the objective function reflects the demand curves of the connections.

Along implementation-based lines, a promising direction would be to deploy the allocation policy we have described here in existing networks which support rate-based allocation. An important aspect of deployment would be to expand on the heuristics we describe for dynamic adjustment of the allocation as bulk connections arrive and depart, and to define the integrated service mechanisms needed to allow bulk connections to coexist with interactive and streaming applications. In practice, it might also be necessary to guarantee some minimal level of service to bulk transfers which are present but do not contribute to the throughput maximizing allocation.

For data selection, deploying the policies in existing network applications is considerably easier, as it can be done at the application level, without modifying existing protocols. The challenge behind the implementation of reliable bulk multicast with forward error correction is to judiciously choose the right settings for the block size and the amount of redundancy. Ideally, one would hope to keep W relatively small to minimize the delay of reconstructing each message block and to keep c large, but small enough to tolerate a bursty period of high packet loss. This can be complicated by the fact that the optimal values for the settings may change over time. Another research possibility for reliable multicast is to incorporate forward error correction with other approaches to minimize the number of retransmissions, such as local recovery. Finally, there are several other multicast applications, such as reliable layered multicast and unreliable streaming applications, for which one could quantify reasonable data selection policy performance measures. The design of data selection policies which perform well with respect to these measures would be another interesting research direction.