Chapter 4

TCP Performance over Satellite Links

The TCP/IP protocol suite has become the lingua fraca of data communications, and the TCP protocol is used for most communications that require guaranteed, end-to-end reliability. Unfortunately, the performance of TCP is often degraded when the end-to-end path includes a satellite link. In this chapter, we focus on the problem of improving TCP performance over satellite links.

In recent years, the subject of TCP over satellite links, and more generally over wireless links, has been a fruitful research area. Nevertheless, the performance of modern TCP implementations over satellite links is still disappointing. In the Internet, satellite links are often used in the configuration depicted in Figure 4.1. In this configuration, the satellite provides access to the wired Internet. Furthermore, the cost of satellite transponder access generally dictates using the satellite link in an asymmetric bandwidth configuration (with more bandwidth allocated in the direction to the clients), or even in a hybrid unidirectional (broadcast) configuration with a telephone return path [100]. This kind of configuration has not been treated thoroughly in the literature, and in considering the configuration shown in Figure 4.1, we were led to consider the following research problems:

1. Satellite TCP connections for which a portion of the connection traverses the wired Internet are subject to severe throughput degradation if the packets flow through a queue that is being congested by connections with a short round-trip time (RTT). Can this bias against long RTT connections be overcome by simple changes to the congestion avoidance algorithm in end hosts?

2. In the current Internet, there exists a wide variety of TCP implementations with various options that interact in different ways. What is the best combination of (standard) TCP options and implementation guidelines for use over satellite channels?

3. How much performance advantage can be gained by "splitting" a TCP connection at a gateway located at the satellite terminal equipment connected to the wired Internet, thereby shielding the satellite subnetwork from the rest of the Internet?

4. In the case of split connections, how much further improvement could be gained by using a transport protocol specifically optimized for the satellite environment?

In this chapter, we focus on the first three questions raised above, and defer the fourth to Chapter 5. The first two questions mainly relate to improving various aspects of TCP's intertwined

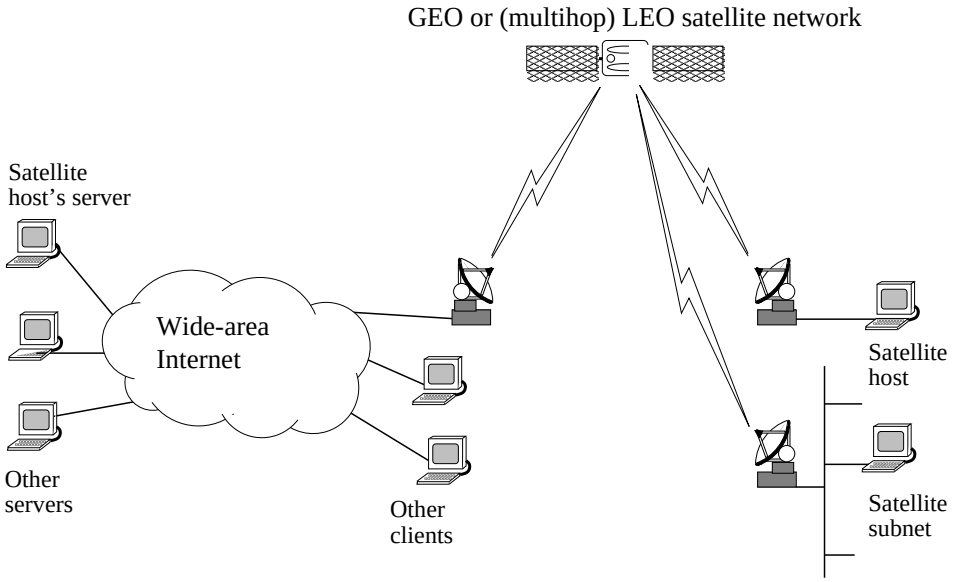

Figure 4.1: Example of a broadband satellite network in which a satellite-based host communicates with a server in the Internet.

congestion control and *loss recovery* mechanisms. To address the first question, we investigate the potential improvement of changing the end-to-end behavior of the congestion avoidance algorithm. Next, we investigate the file transfer behavior of different variants of TCP SACK and present a standards-conformant algorithm that achieves high performance in a satellite environment. Finally, we address the third question by investigating the potential benefit of splitting the end-to-end connection at a gateway.

4.1 TCP Fairness in a Heterogeneous Environment

4.1.1 Introduction

The fairness problem in TCP is rooted in its congestion avoidance mechanism, which we described above in 2.1.1. The congestion avoidance phase is sometimes referred to as "additive increase and multiplicative decrease," because, in the absence of congestion, segments are added to the window over time, while in the presence of congestion, the window is halved (or multiplied by one half).

The "additive increase and multiplicative decrease" algorithm in TCP parallels a similar algorithm in the DECnet protocol [66]. Chiu and Jain showed that this algorithm leads to fair allocations of network bandwidth even though it operates in a distributed manner [28]. However, their analysis presumes that all connections in the network share the same additive increase rate and multiplicative decrease factor. In TCP, the multiplicative decrease factor (1/2) is the same for all connections, but the policy of an additive increase of one segment per round trip time (RTT) does

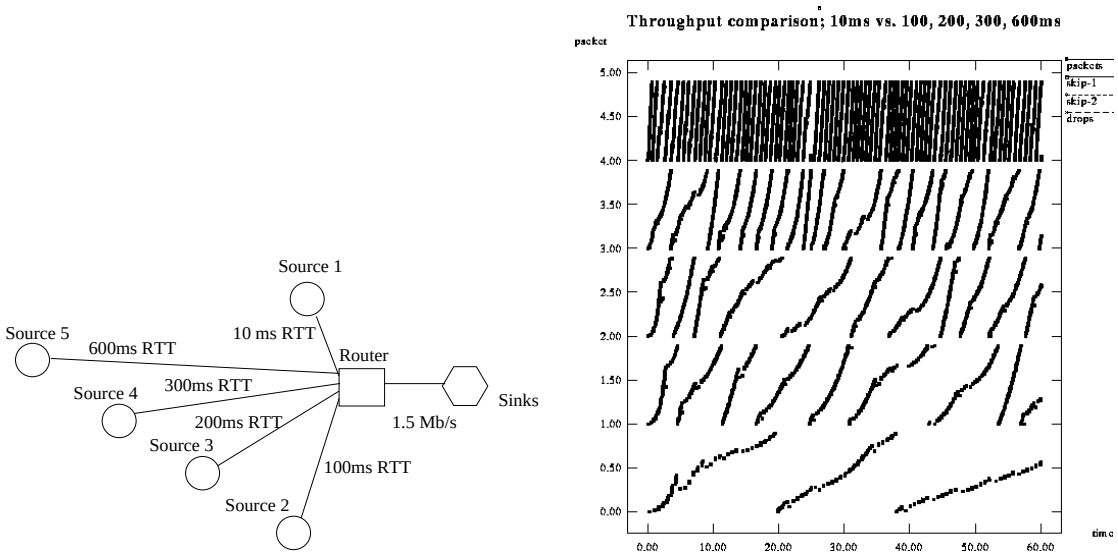

(a) Congested network topology. (b) TCP sequence number evolution vs. time.

Figure 4.2: A demonstration of the unfairness of the current TCP congestion avoidance algorithm. The connections, from top to bottom, have RTTs of 10, 100, 200, 300, and 600 ms respectively.

not provide a uniform increase in the rates of TCP connections with different RTTs[1]. In particular, connections with long RTTs open their window more slowly than those with short RTTs. And if a mixture of such short and long RTT connections share a bottleneck link, severe unfairness is inevitable as the short RTT connections grab the available bandwidth well before the long RTT connections have a chance [54].

Figure 4.2 illustrates an example of this problem by showing simulation results for a 60 second trace of TCP connections over the illustrated topology. In Figure 4.2, the evolution of the sequence number is plotted for 5 connections (from top to bottom, with RTTs of 10, 100, 200, 300, and 600 ms, respectively) sharing the same bottleneck link. The sequence number in this simulation is on a per segment basis, and the plots wrap after every 90 segments. The long RTT connections do not obtain an allocation close to their fair share of the bottleneck link, and their overall throughput performance suffers drastically.

To combat the bandwidth inequities that result from heterogeneous RTTs, Floyd proposed a modification to TCP's window adjustment algorithm that counteracts the RTT bias. In this section, we elaborate Floyd's "Constant-Rate" algorithm [39] with a thorough investigation of the performance achievable by both universally and selectively (i.e., incrementally) deploying a TCP with a modified window increase policy in the congestion avoidance phase of the connection.

Floyd [39] developed a fairly general characterization of window increase algorithms that facilitates the discussion of fairness. Although TCP maintains its send window in units of bytes, we find it more convenient herein to discuss it in units of segments. A key assumption is that a number of segments approximately equal to the send window size is sent every RTT; this is generally true for long RTT connections. Let c be the increase (in segments) in the size of the send window for a

[1] In the following, we distinguish between the overall congestion avoidance $algorithm$ and the $policies$ implemented in this algorithm such as the multiplicative decrease factor of 1/2 and the additive increase of one segment per RTT.

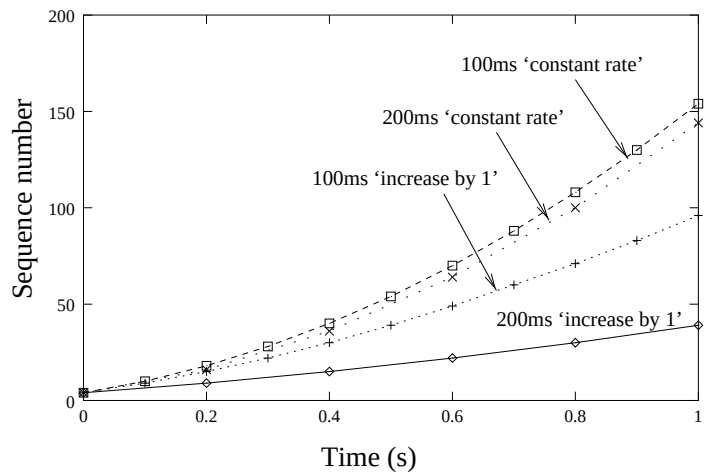

Figure 4.3: Simulated sequence number evolution for connections with different RTTs under both the constant rate and increase-by-one policies.

connection in one round trip time RTT. Therefore the window grows at a rate of c/RTT segments per second when additive increase is in effect. In conventional TCP, $c = 1$. Floyd refers to the standard TCP policy as an "increase by 1" policy.

If one were to scale the window growth rate c/RTT by RTT, the effect would be to build the window at a constant rate of c segments per second, independent of the RTT. However, the window growth rate does not equal the growth rate in data transmission. As [41] points out, it results in a "linear-in-RTT" bias in the sending rate. Because each connection can send a full window's worth of segments each RTT, shorter RTT connections achieve greater throughput over a common time interval. To fully remove the bias, we must change the additive increase to $c * RTT$ segments per second; i.e., a factor of RTT^2 faster than the original algorithm. Floyd defines such an increase as a *Constant-Rate (CR) increase policy*, since it can roughly be interpreted as causing the rate of segment transmission to increase at a constant rate.

Figure 4.3 demonstrates the behavior of two of these policies for two connections with different round trip times. The figure plots the equation

$$P_{c_i,RTT}(t) = \sum_{k=0}^{\lfloor t/RTT \rfloor} (cwnd + kc_i),$$

for the different RTTs and policies c_i. This equation describes the number of segments sent ($P_{c_i,RTT}(t)$) assuming a window of segments is sent each RTT and there are no losses, where c_0 is equal to 1 (the standard policy), and c_1 implements the constant rate policy (and hence is a different value for each RTT). While only a rough approximation, the graph confirms the shape of the increase rate curve for each policy.

4.1.2 Methodology

We used the UCB/LBNL Network Simulator "ns" [2] to evaluate TCP performance. In addition to using the standard ns modules, we ported the HTTP traffic generator module from Bruce Mah's INSANE simulator[3], which gave us the ability to add a mix of realistic background traffic to our simulations.

Our study focused on large file transfer performance. While short HTTP transfers and Telnet connections over long RTT paths are also subject to performance degradation during periods of congestion, this degradation is due more to the fundamental latency of long RTT connections than to problems with congestion avoidance. Additionally, HTTP protocol implementations are migrating towards "persistent-HTTP" and longer duration TCP connections. We also did not assume the implementation of fair scheduling and TCP-friendly buffer management that can isolate flows or classes of flows from one another (e.g., as discussed in [131]), or pricing structures that might give network providers incentives to protect the throughput of paying customers. In short, we assumed an environment similar to the present day Internet, with the addition of Random Early Detection (RED) queues [44], and the latest in standardized TCP improvements (Selective Acknowledgments (SACK) [83] and large window enhancements [67]).

Performance Metrics

A number of metrics for quantifying fairness have been proposed but no single metric has common acceptance [39]. In this paper, we consider a "fair share per link" metric; i.e., if there are n flows through a bottleneck link, each flow has the right to $1/n$th of the capacity of that bottleneck link. Jain's metric of fairness [28] is applicable in this context. For n flows, with flow i receiving a fraction b_i on a given link, the fairness of the allocation is defined as:

$$Fairness \equiv \frac{(\sum_{i=1}^{n} b_i)^2}{n * (\sum_{i=1}^{n} b_i^2)}.$$

This metric ranges continuously in value from $1/n$ to 1, with 1 corresponding to equal allocation for all users. Utilization is another important metric, since high fairness is of little use if the link capacity is grossly underutilized. Utilization is defined herein as the number of original bits (i.e., not counting retransmissions) successfully transferred over a link during some time interval divided by the product of link rate and that time interval; this is often called "goodput."

$$Utilization \equiv \frac{(segments_acked) * (segment_size)}{(link_rate) * time}.$$

Topologies

We explored a number of test configurations that allowed us to isolate selected behavior of both the standard and our proposed window adjustment policies. We selected the topologies illustrated in Figure 4.4; similar test configurations have previously been used by the research community to study the effects of congestion. The first test configuration (Topology 1) was used to

[2] http://www-mash.cs.berkeley.edu/ns/

[3] http://http.cs.berkeley.edu/∼bmah/Software/HttpModel/

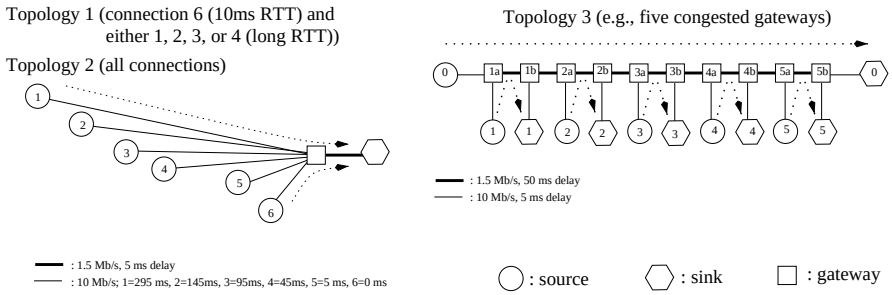

Figure 4.4: Three simulation topologies.

study the effects of two long duration connections, one with a short RTT (10 ms) and one with a long RTT (100-600 ms), sharing a single bottleneck link. The second test configuration (Topology 2) was used to examine the effects of many competing connections over a single bottleneck link. The six connections have RTTs of 10, 20, 100, 200, 300, and 600 ms. The third configuration (Topology 3) was used to examine the effects of long RTT connections that must also traverse a number of network hops populated by short RTT connections. This topology is very similar to the one previously used by Floyd to study the CR algorithm [39]; the number of congested gateways could vary between 1 and 10 (the figure illustrates 5 congested gateways).

Configuration Details

We studied the performance of two different TCP variants described above in Section 2.1.1: TCP NewReno, and TCP SACK. We note that other researchers have detected problems with using TCP Reno (a version of TCP that does not perform adequately when multiple drops occur in a window of data) in combination with congestion avoidance mechanisms that try to add more than one segment per RTT [17]; therefore, we avoided such implementations. We also examined two different queueing schemes: traditional "first-in, first-out" (FIFO) queueing, and Random Early Detection (RED) with packet discard.[4] In our simulations, data packet sizes were fixed at 1000 bytes, and the bottleneck link speed was 1.5 Mb/s. We examined a range of queue sizes from 4 to 50 packets, but in the data that follows, we concentrate on a RED queue size of 50 packets with a "minimum threshold" of 20 packets and a "maximum threshold" of 40 packets; all other RED parameters were set to the ns defaults. 20 packets in this case is approximately 100 ms at our output line rate.

[4]RED queues operate by computing an exponentially weighted moving average of the queue size. When the average queue size is below some minimum threshold, the queue does not drop any packets. When the average queue size is between the minimum and maximum threshold, the queue probabilistically drops incoming packets according to an algorithm described in [44]. When the average queue size exceeds the maximum threshold, the queue drops every incoming packet. The instantaneous queue depth can exceed the maximum threshold if the average queue depth is below the maximum threshold.

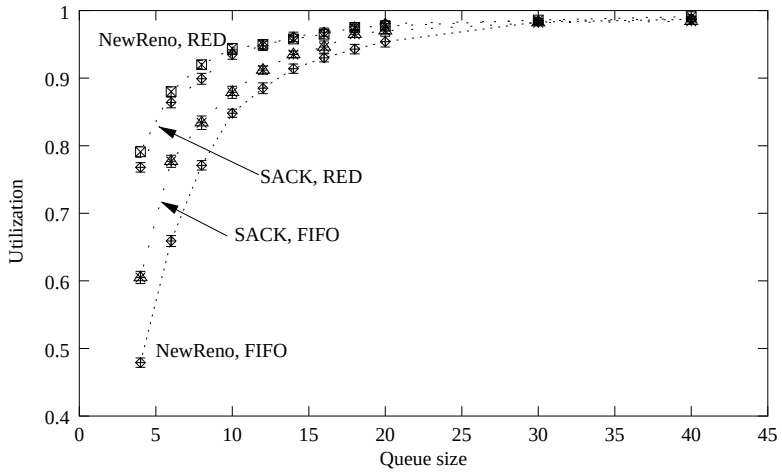

Figure 4.5: Benchmark performance results using Topology 2: Utilization vs. queue size.

Data Analysis and Presentation

TCP throughput in an environment containing random traffic can be quite variable, because small changes in initial conditions can cause wide variations in resulting behavior. Therefore, we computed the utilization and fairness of a particular configuration as follows. We first ran enough independent simulations such that the sample standard deviation of each connection's throughput was within 5% of its sample mean (this generally required around fifty runs). We then used these sample means to compute the fairness and utilization of a given topology. In the remainder of Section 4.1, if the experimental data does not explicitly list error bars, the reader may assume that the sample standard deviation is within 5% of the value listed. In the following subsections, we first provide some benchmark data, followed by an analysis of the Constant-Rate policy, followed by experiments aimed at selectively increasing the aggressiveness of a long-delay TCP connection.

4.1.3 Benchmark Results

To calibrate our simulation studies, we established a set of benchmark performance results for each of our test topologies. Though used principally to gauge the efficacy of our proposed policies, the benchmark data itself reveals some interesting effects. In this section, we examine benchmark performance data from Topology 2 in Figure 4.4.

Figures 4.5 and 4.6 plot the utilization and fairness of the standard TCP window adjustment policy for Topology 2. The error bars on these figures represent 99% confidence intervals, and are very small. We show results from the combination of two TCP variants, NewReno and SACK, with two queueing disciplines, FIFO and RED. In this representative data set (and in our other benchmark data), the following trends are evident:

- The utilization of the bottleneck link improves with increased queue size, because, under congestion, large queues keep the link busy as they drain out while TCP sources back off

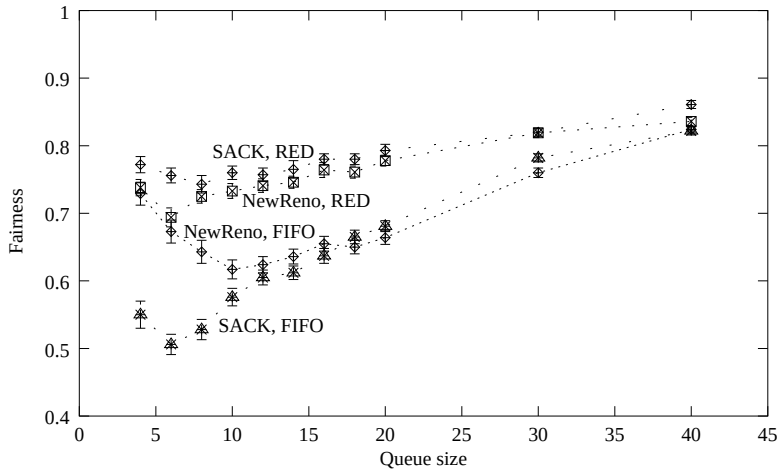

Figure 4.6: Benchmark performance results using Topology 2: Fairness vs. queue size.

their sending window. Additionally, larger queues absorb bursts of packets and prevent coarse timeouts, which cause long idle periods.

- Network fairness is poor in almost all cases, except when queue sizes are very large. In general, the two short delay connections obtained roughly 50% of the bandwidth, and the background WWW traffic consumed 20% of the bandwidth. The remaining 30% was split unequally among the 4 long RTT connections, with the longest connection receiving only about 24 to 64 kb/s (2 to 4%) for TCP NewReno with FIFO queueing. While even larger queue sizes may help further, they would also introduce more significant delay variability.

- In general, when queue sizes are reasonably large and when all TCPs use SACK instead of NewReno, the network fairness is marginally better. This is most likely due to SACK's superiority in recovering from multiple drops in a single window. Since multiple congestive losses in a single window are more likely to occur in a connection with a long RTT, the use of SACK helps such connections. However, the use of SACK by all TCP connections does not, by itself, remove the bias against long RTT connections.

In general, RED queues perform much better in terms of utilization and fairness than do FIFO queues. However, we found that the use of RED and SACK alone, without modifications to TCP sending behavior, still leaves much room for improvement in fair bandwidth allocation. RED queues equalize the bandwidth of flows with similar RTTs, but do not do so for flows with heterogeneous RTTs, as pointed out in [44]. In the remainder of this section, we focus on enhancing the performance of TCP implementations that use SACK and networks using RED queues.

4.1.4 Performance of the Constant-Rate Policy

In this section, we describe the case in which each foreground and background TCP connection uses a Constant-Rate (CR) policy. In TCP implementations, an additive increase to the TCP

variable snd_cwnd (send congestion window) of approximately one segment per RTT, assuming an acknowledgment (ACK) is received for each segment, is effected by executing the following pseudocode upon receipt of a new ACK:

```
snd_cwnd = snd_cwnd + 1/snd_cwnd.
```

In this manner, the TCP connection gradually adds to its congestion window at the rate of approximately one segment per RTT; this approach to building the congestion window reduces transmission burstiness [66]. To implement a CR policy, we can modify the window increase algorithm to account for the RTT bias:

```
snd_cwnd = snd_cwnd + (c*rtt*rtt)/snd_cwnd,
```

where c is the constant that controls the rate. This policy causes an additive increase in the throughput rate that is the same for all connections. After initial experiments, we observed that the second term of the above equation could lead to very bursty send patterns, which led to increased losses. For example, if the RTT is large and the value of snd_cwnd is small, each ACK can trigger the transmission of several segments. To avoid this behavior, we bounded the increase per ACK by 1 segment; i.e.:

```
snd_cwnd = snd_cwnd +
min((c*rtt*rtt)/snd_cwnd, 1 segment).
```

With this constraint on the sender's behavior, the TCP connection is never more bursty than a TCP connection in slow start. Another approach, with which we did not experiment, would be to smooth the sending of several segments across a longer time period.

One question previously raised by Floyd is how to pick the proper value for the constant c. One way to think of the value of c for CR connections is how the aggressiveness of the CR connection would compare to that of a standard TCP connection with a certain RTT. For example, if $c = 100$, the value of the RTT that makes the numerator equal to 1 in our pseudocode above is 100 ms. Therefore, an environment in which $c = 100$ would have connections that were about as aggressive as normal TCP connections with 100 ms RTTs. We chose to experiment with a range of values, between $c = 4$ (as aggressive as standard 500 ms connections) and $c = 1600$ (25 ms).

In addition to varying the constant c, we experimented with several other variations in an effort to identify which types of environments were suitable for the CR policy:

- TCP NewReno vs. TCP SACK,
- RED vs. FIFO gateways,
- bottleneck queue lengths (or RED maximum queue thresholds) from 4 to 50 packets, and
- TCP RTT timer granularity of 500 ms (standard in many TCP implementations) vs. 10 ms.

As mentioned above, Topology 3 conforms closely to one with which Floyd experimented [39], and we were able, when using a similar value for the CR constant ($c = 4$), to confirm their results that CR can substantially improve the fairness of connections traversing multiple gateways when all connections use a CR policy. However, we also observed the following general trends in our data:

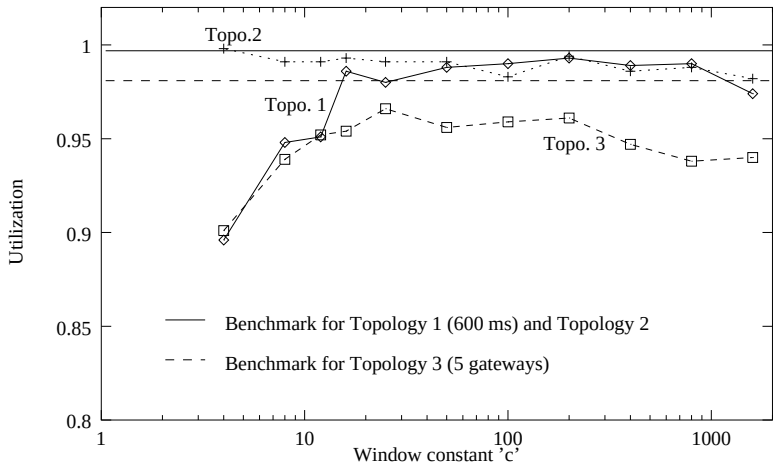

Figure 4.7: Utilization vs. window constant of TCP SACK with fine-grained RTT estimates over topologies with bottleneck RED queues.

- Deep RED queues appear to be a prerequisite for good performance of the CR policy. Performance when FIFO queues or short queues were used was very inconsistent, in the sense that there was often no value of c that simultaneously yielded high fairness and high utilization. Moreover, we could not determine strong correlations between the value of c and the fairness and utilization metrics we were using; i.e., the performance was highly sensitive to the particular simulation topology.

- TCP SACK and fine-grained RTT timers were the next most important indicators of good CR performance. The use of SACK helps TCP recover from losses more quickly, which leads to improved and more consistent performance. Also, many existing TCP implementations use a coarse estimate of the RTT, which impairs the ability of our modified congestion avoidance algorithm to determine the true RTT of the connection.

Figures 4.7 and 4.8 plot the utilization and fairness performance of TCP SACK over bottleneck RED queues when all connections, including background HTTP traffic, use the same CR policy, constant c, and RTT timer granularity of 10 ms. Topology 1 corresponds to the case in which the long RTT connection has a round trip propagation delay of 600 ms, while Topology 3 in this case corresponds to the topology with 5 congested gateways (the trace is taken from the first congested gateway). For comparison, we also plot as horizontal lines the utilization and fairness achieved when all TCP connections use the standard algorithm (i.e., benchmarks). The data indicates that the fairness can be substantially improved if all connections adopt the CR policy. However, the utilization suffered for small c when there were only two foreground connections in the topology (Topologies 1 and 3). When statistical multiplexing was in full effect (Topology 2), both fairness and utilization were near optimal for small values of c. Additionally, in Topology 3, although the fairness improved substantially, equal allocations were not obtained by the CR policy because the long RTT connection is also traversing multiple congested gateways.

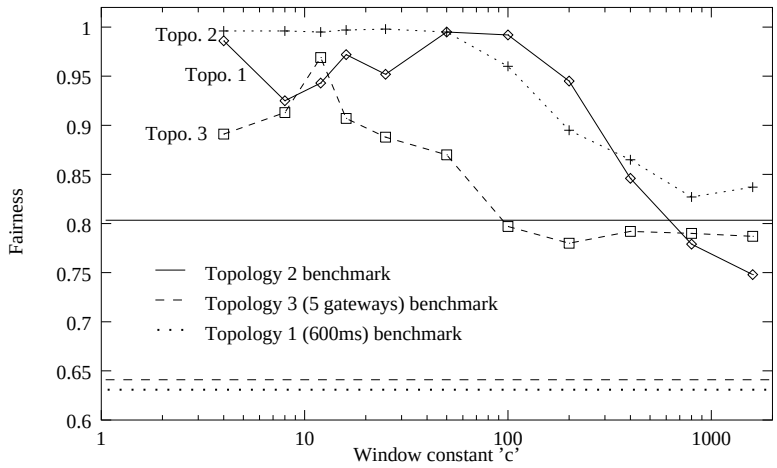

Figure 4.8: Fairness vs. window constant of TCP SACK with fine-grained RTT estimates over topologies with bottleneck RED queues.

In these experiments, as c became larger, connections became more aggressive, to the point that the bound in our policy of adding 1 segment per ACK was in effect nearly all of the time. Consequently, since the CR policy was no longer being applied, the unfairness reappeared. In general, we observed that the fairness properties were best when the value of c was below 100. However, if too few connections are using the link, such as in Topologies 1 and 3, such a small value of c can lead to lower utilization. Because it is difficult in practice for a given connection to determine the number and type of connections against which it is competing, we conclude the following negative result: a good choice of the constant c cannot be determined with high confidence on an operational basis.

Not only does the CR policy appear difficult to manage in a distributed network, we also found it susceptible to the presence of TCP connections operating under the standard policy. For example, Figure 4.9 illustrates the fairness performance when a single additional connection using the standard window increase policy was introduced into each of the topologies (and also included in the fairness computation). Although this additional connection slightly improved link utilization, much of the fairness improvement due to CR was lost when this competing connection was introduced, as it not only used a disproportionate share of the bandwidth itself but also acted as a "trailblazer," improving the performance of short RTT connections that were using the CR policy by a disproportionate amount. We also observed similar performance degradation if no extra file transfers were introduced, but instead the HTTP background traffic (20% of the bottleneck link rate, on average) used the standard policy. Similar effects (passive connections competing with more aggressive connections using the standard congestion avoidance algorithm) are also responsible for the poor performance of TCP "Vegas" in a long RTT environment characterized by a mix of heterogeneous TCP implementations [147].

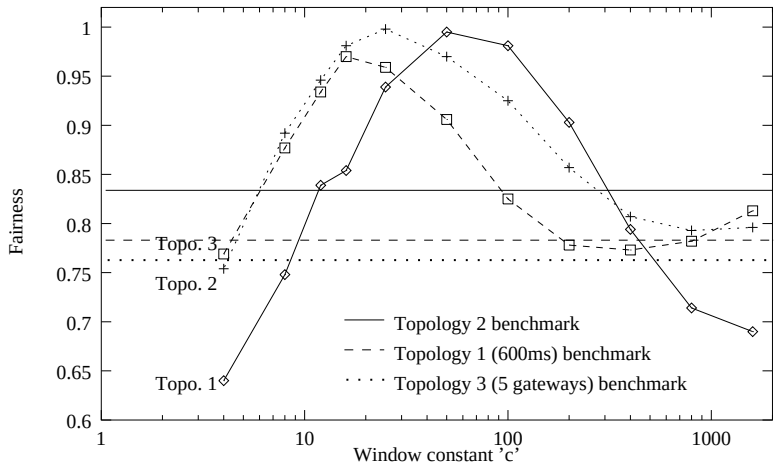

Figure 4.9: The sensitivity of the proper selection of the constant c when the effects of a standard TCP connection are added to the performance shown in Figure 4.8.

4.1.5 Selectively Modifying the Additive Increase Policy

We next investigated whether the throughput of an individual long RTT connection could be improved by modifying the additive increase policy of only the long RTT connection. We were interested in two questions:

• Can an individual connection improve its own throughput by becoming more aggressive during additive increase?

• If so, how does the individual connection's more aggressive behavior affect the performance of other (unmodified) connections using the same path?

To study the first question, we experimented with an "increase-by-K" (IBK) policy rather than the standard "increase-by-one," again limited by a maximum increase of one segment per ACK. In other words, we used the following pseudocode in our implementation:

```
snd_cwnd = snd_cwnd +
min((K/snd_cwnd),1 segment).
```

For example, by setting $K = 2$, we built the window by roughly 2 segments per RTT.[5] Again, the increase is bounded by 1 segment per ACK, but this is the general trend.

Figure 4.10 illustrates fairness results, calculated over the entire simulated network, for the case in which only one connection in the simulated topology used the IBK policy. In particular, we enabled the IBK policy on the longest RTT connections in each of the three topologies, and then repeated the experiment by enabling the IBK policy on only the 300 ms connection in Topology 2. In the graph, the value of $K = 1$ (left-most data points) corresponds to the normal (benchmark) case. We observed that the long RTT connection was able to steadily improve its performance over that of the benchmark case by increasing K across the range of values we considered. This resulted

[5] If delayed acknowledgments are being used, this is akin to correcting the window growth penalty that is due to delayed acknowledgments.

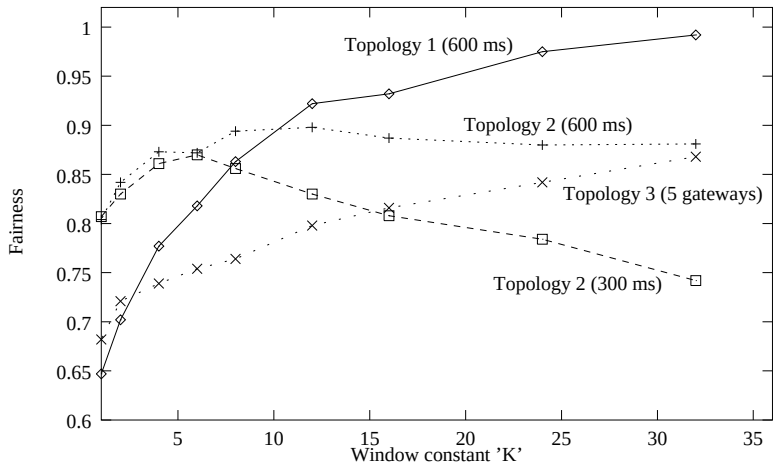

Figure 4.10: Improvement in fairness vs. window constant due to the IBK policy.

in improved fairness in all topologies for small values of $K > 1$. For larger values, even though the throughput of the long RTT connection continued to improve, fairness actually decreased in some topologies as the more aggressive connection began to take more bandwidth than its fair share.

In Figure 4.10 we plotted the performance of TCP SACK over RED queues, and found that there was no limit to the improvement that a more aggressive connection could obtain for itself. We repeated the experiment for TCP NewReno over FIFO queues, and found that connections could increase their own performance, independent of the topology, by using a value of K of up to 4 or so. However, for higher values of K, performance degraded, because the sending behavior became too bursty for the FIFO queues to successfully absorb.

Because the performance improvements result from increasing TCP's aggressiveness, we should be concerned that this can have a negative impact on other peer connections. Remarkably, we found that, in every case we examined, the average fairness index *always* improved, and the average utilization held relatively constant, when the more aggressive connection used a modest value of K (less than 8 or so). This improvement occurred regardless of whether TCP SACK or NewReno was used, or whether FIFO or RED queues were present. In fact, the majority of the redistributed bandwidth came from connections that were already using more than their fair share. The effect on other connections was similar to what they would have experienced had the long RTT connection actually been a connection with a somewhat shorter RTT.

For example, Table 4.1 provides an example of the magnitude of the performance gains achievable. For the value $K = 4$, we tabulate the utilization, fairness, and throughput of the four longest RTT connections in Topology 2. In the first column are results from when each connection used the standard policy, the second column shows the results from when only the 300 ms connection used an IBK policy, and the third column is from when only the 600 ms connection used an IBK policy. This table illustrates that the more aggressive connections did not seriously harm the throughput of the peer connections. The throughput gains (highlighted in bold font) are substantial; in many cases, throughput increases (and hence reductions in user-perceived latency) exceeded

	benchmark	**300 ms**	**600 ms**
Utilization	0.99	0.99	0.99
Fairness	0.807	0.861	0.873
	Throughput (kb/s)		
600 ms	**68.1 (0.9)**	59.6 (0.6)	**138.7 (2.0)**
300 ms	**116.2 (1.1)**	**240.8 (2.8)**	109.1 (1.1)
200 ms	156.6 (1.7)	139.8 (1.5)	151.6 (1.1)
100 ms	217.1 (2.1)	188.1 (1.9)	207.0 (1.8)

Table 4.1: Effect of the IBK policy on throughput (Topology 2, K = 4). 99% confidence intervals are shown in parentheses.

100% when RED queues were used, and 50% when FIFO queues were traversed.

We next investigated whether the performance gains were sustainable when multiple connections become more aggressive by examining this case with Topology 2. We found that the aggressive connections were able to simultaneously improve their own performance, although their relative performance gains were not what they would have achieved had they been the only aggressive connection. Of course, if every connection adopted an IBK policy, the fairness situation would be back to the standard case, so there must be some RTT threshold beyond which connections can become more aggressive if such a policy is to work in practice. Finally, we experimented with the case in which congestion was induced in both directions of data transfer. This had the effect of disrupting the ACK stream to some extent, but did not significantly affect our main results.

4.1.6 Implementation Issues

We have already discussed some minor implementation changes to the code segment which builds the congestion window. Throughout the discussion, we implicitly assumed that the TCP connection had an accurate estimate of its RTT. In practice, this is not the case. TCP does maintain a smoothed round trip time ($srtt$), but because of the timer granularity of 500 ms in TCP, this value is not very accurate. It is, however, rather easy to improve the RTT accuracy through use of the TCP timestamps option [67]. A sending TCP implementation can put a more accurate timestamp in the TCP timestamp field, which is merely reflected by the receiver; such a technique is suggested for TCP Vegas [16]. However, it is important not to base the retransmission timer value on this accurate timestamp, because TCP fast retransmit and fast recovery rely on the $srtt$ variable being somewhat larger than the actual RTT. We experimented in ns with running the TCP timer granularity at 1ms instead of 500 ms, and found that the more accurate $srtt$ value caused coarse timeouts to trigger before fast recovery could be accomplished.

One practical issue is that modifications to the sending algorithm of an implementation have little use if the implementation is a client of a large data transfer rather than the source of the data. However, it is possible to indirectly modify the sender's behavior by actions taken at the receiver. For example, by sending more ACKs back to the sender, the receiver can "speed up" the sender; e.g., if one were to receive a segment of 1000 bytes, sending ACKs 1:250, 251:500, etc. would quadruple the window build rate. We did not experiment with this technique, and note that it

is of limited utility when the reverse channel is constrained and cannot handle the additional ACK traffic.

Another alternative to increasing the aggressiveness of a single connection is to run multiple connections in parallel, coordinated by the application or some type of session manager. This technique, sometimes called "striping", has been in use for some time by satellite operators and WWW browser software. One advantage of this approach is that it overcomes small offered windows by the receiver. This technique, however, can impact the network more than our approach since multiple slow starts are launched into the network simultaneously. Recent results in which congestion window state is shared across the multiple connections can potentially combat the problem [11]. In general, the use of striping without the constraints outlined in [11] are not favorably viewed by the research community.

4.1.7 Summary

In this section, we have presented the results of our investigation of simple changes to TCP's congestion avoidance algorithm in an effort to improve its fairness properties. While we found that the Constant-Rate (CR) policy could improve fairness dramatically, we faced two practical difficulties that would likely prevent universal deployment of this scheme in its current form: i) the proper selection of a constant is dependent upon the network topology and the number of peer connections and is therefore difficult to determine in a distributed manner, and ii) the fairness benefits of the CR policy can be confounded by competing connections using standard congestion avoidance, thereby making it disadvantageous to deploy CR in an existing heterogeneous environment. However, when we instead made only certain long RTT connections slightly more aggressive, we were always able to improve network fairness while keeping bottleneck link utilization relatively constant by using an increase-by-K (IBK) policy. Interestingly, the effects on other unmodified connections that were sharing the bottleneck link were similar to what they would have experienced had the modified connection actually been a connection with a shorter RTT.

Our results indicate that it may be beneficial for long RTT connections (running TCP SACK) to become slightly more aggressive during the additive increase phase of congestion avoidance. While our data set is not comprehensive enough to allow us to advocate a particular policy at this time, the IBK policy for small values of K (such as 2 or 4) may significantly improve the throughput while not significantly impacting other flows. Such a policy could be invoked in practice when the TCP implementation detects that the connection has an RTT above a certain threshold (for example, connections traversing a GEO satellite link have a much larger RTT– at least 500 ms– than terrestrial connections). A TCP receiver could even induce the sender into an IBK policy by acknowledging data in smaller chunks. Determining appropriate values for K as a function of RTT, as well as determining the accuracy and resolution required of TCP's RTT estimates, could be the focus of future work.

4.2 End-to-End TCP Performance over Satellite Links

4.2.1 Introduction

In this section, we quantify just how well state-of-the-art TCP implementations perform in a satellite environment composed of one or more satellites in geostationary orbit (GEO) or low-

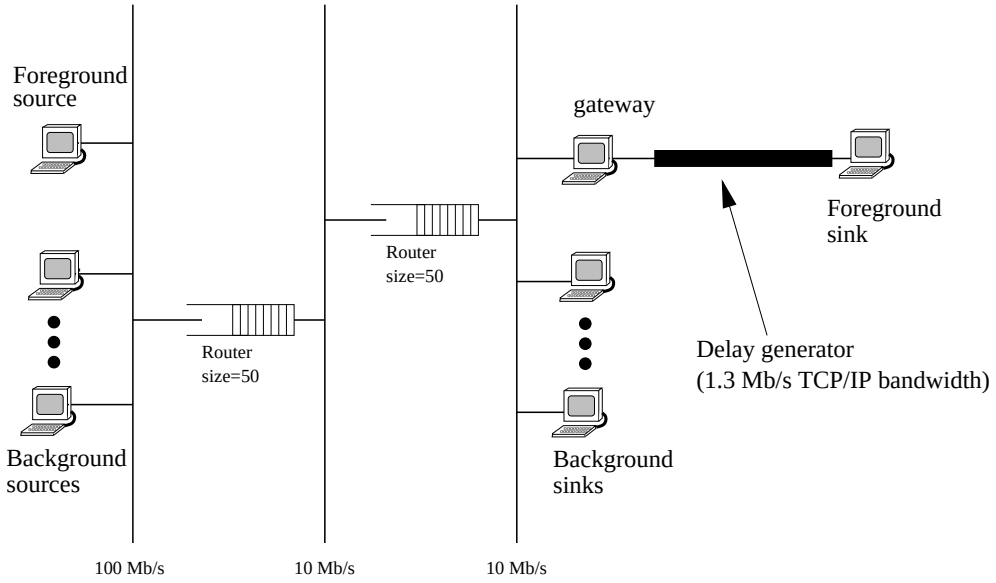

Figure 4.11: Configuration for network experiments.

earth-orbit (LEO), particularly when the satellite connection forms only a part of the end-to-end connection, as shown in Figure 4.1. We focused on two types of workload found most commonly in the Internet: large file transfers, and short Web connections.

Our assumptions about future satellite network characteristics are shaped by projections of future commercial systems (e.g., Teledesic [130], Spaceway [38]) that will offer Internet connections at up to broadband (tens of Mb/s) data rates via networks of LEO or GEO satellites (or hybrid constellations). Users may contact other hosts in either the satellite network or the wide-area Internet. We discussed some of our assumptions about the transmission and congestion characteristics of the end-to-end path using such satellite systems in Section 2.1.2. In short, we assume future satellite networks characterized by low BERs, potentially high degrees of bandwidth and path asymmetry, high propagation delays (especially for GEO based links), and low internal network congestion. These assumptions were used to drive our protocol design and performance analyses described in the rest of this chapter.

4.2.2 Methodology

This experimental methodology pertains to the remainder of this chapter and also to the results presented in Chapter 5.

Experimental Setup

Our experiments were conducted using hosts, running BSD/OS 3.0 UNIX, connected to Ethernets in a local-area subnet at Berkeley. The TCP implementations on these machines are derived from 4.4BSD-Lite (also known as Net/3 [144]), with modifications to support our experiments.

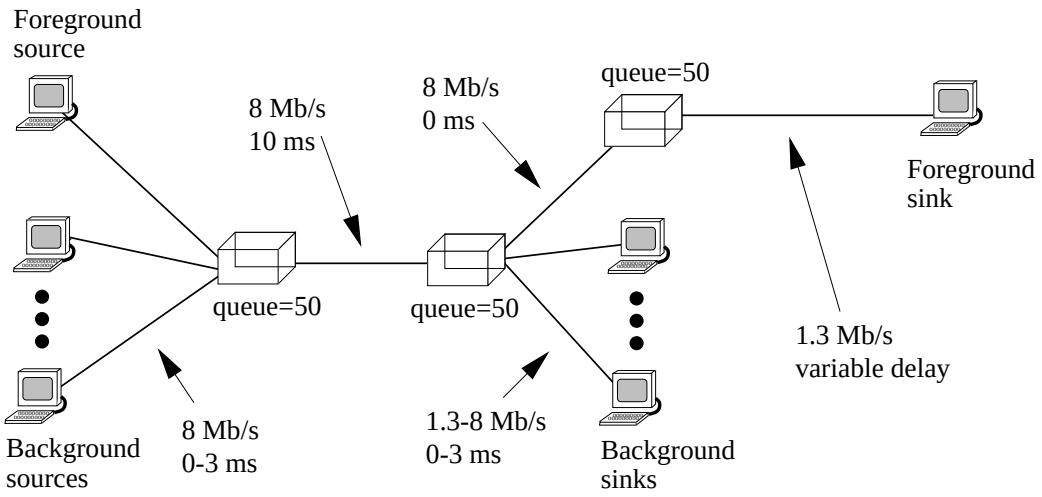

Figure 4.12: Configuration for simulation experiments.

We configured the receivers to offer the largest window possible (240 KB) to the senders. For the experiments, traffic sources were connected to a 100 Mb/s Ethernet, and traffic sinks were on a 10 Mb/s Ethernet separated by a 10Mb/s transit Ethernet segment. Figure 4.11 illustrates the experimental topology. To generate traffic, we used a combination of the `sock` program [127] for bulk file transfers and a HTTP traffic generator for testing of "4K slow start" and T/TCP. This traffic generator generated small file transfers according to empirical distributions drawn from Bruce Mah's HTTP traces [79]. We implemented STP in the BSD/OS UNIX kernel.

For investigating satellite transport protocol performance, it is usually sufficient to experiment with delay and error simulators rather than with detailed emulators of the transmission channel. To emulate satellite links, we used modified device drivers that delayed sending a packet onto the Ethernet for a deterministic amount of time. These drivers can also constrain the maximum rate at which a host can send data. We modeled GEO satellite links by a constraint of 1.3 Mb/s of TCP/IP bandwidth (i.e., approximately T1 rate at the physical layer), on a 600 ms RTT link. LEO satellites were modeled by a constraint of 1.3 Mb/s with a fixed RTT in the range of 40-400 ms [49]. Our links had no bit errors or variation in propagation delay, which, while not representative of all satellite links, exemplifies the common case.

In addition to controlled experiments performed in our local environment, we also describe experiments in Chapter 5 involving two commercial networks in our wireless networking testbed. We used a network based on a direct broadcast satellite (the Hughes DirecPC system, which covers the contintenal US), and a packet radio network (the Metricom Ricochet system, deployed in the San Francisco Bay area). For DirecPC experiments, we sent data from a computer located at the DirecPC uplink center at Germantown, MD over the satellite link to a multi-homed host on one of our subnets. We used the wide-area Internet to return acknowledgments to the traffic source. To emulate a normal user experience with the DirecPC system, we constrained the return link to be bandwidth limited to 50 Kb/s to simulate a modem connection. Although not a satellite network, the Ricochet network offers a challenging environment for transport connections, includ-

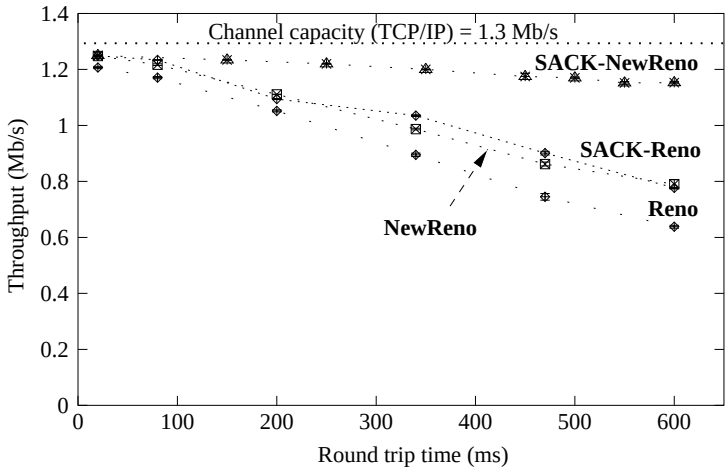

Figure 4.13: Throughput performance of TCP SACK NewReno, TCP SACK Reno, TCP NewReno, and TCP Reno over an experimental path with a TCP/IP bandwidth of 1.3 Mb/s and no transmission errors. Data points represent the sample means from 20 independent transfers of 10 MB each. In this and subsequent figures representing a large number of experimental results, error bars represent 95% confidence intervals.

ing asymmetry and large latencies; we used this network only for testing of the STP protocol as described in Chapter 5. In these experiments, a wired host at Berkeley communicated with a host on the Ricochet network using the packet radio network in both directions.

Simulation Configuration

We used *ns*, described above in Section 3.2, to test simulated topologies that matched our experimental setup. We aligned the TCP modules to match our implementations, and wrote a STP simulation module to closely emulate the implementation used in the experiments. We also used a background HTTP traffic generator, similar to that used in the experiments, to lightly load the network topology and to break up any TCP phase effects [43]. Our simulation topology, which conformed closely to the experimental setup, is shown in Figure 4.12.

4.2.3 Performance for Large File Transfers

TCP is the dominant protocol for file transfers (FTP) in the wide-area Internet. In this section, we describe simulations and experiments used for characterizing file transfer performance over satellite links.

To maintain high throughput for large file transfers, the TCP congestion window must be large. This implies that the congestion avoidance and loss recovery mechanisms are very important in determining performance. In this section we examine the performance of four variants of TCP loss recovery and congestion control, which we first introduced in Section 2.1.1:

- **TCP Reno** The unmodified TCP implementation in our BSD/OS 3.0 operating system is commonly known as TCP Reno. Many modern TCP implementations are largely based on this version of TCP. Of the satellite-friendly TCP extensions described above, BSD/OS 3.0 supports window scale and path MTU discovery.

- **TCP NewReno** TCP "NewReno" is a collection of bug fixes and refinements for how TCP Reno handles the fast recovery phase of congestion avoidance. Our TCP NewReno implementation is similar to TCP Reno except that it avoids false fast retransmissions [60], multiple window reductions in one window of data [36], and constrains the burstiness of the sender upon leaving fast recovery [36]. Specifically, it implements the "Less Careful, Slow-but-Steady" variant of NewReno described in [42].

- **TCP SACK-Reno** Reno congestion avoidance algorithms may be combined with the SACK option for loss recovery to form TCP "SACK-Reno."

- **TCP SACK-NewReno** Likewise, this corresponds to TCP NewReno congestion avoidance with the SACK option for loss recovery.

It is important to emphasize that all of the above implementations would be regarded as conformant to the TCP standards; in practice, many more variants of TCP exist.

For our file transfer experiments, we repeatedly transferred 10 MB files across our testbed while varying the latency of the emulated satellite channel. The file transfers lasted at least 60 seconds, allowing the low throughput of the initial slow start phase to be amortized across the lifetime of the connection. In the simulations, we added a number background HTTP traffic generators to the topology so as to introduce low levels of cross traffic (approximately 10% of the forward throughput of the channel). These traffic generators did not by themselves congest the forward path; the TCP losses were periodically self-induced by the greedy nature of the congestion avoidance mechanism of the persistent file transfers. In the experiments, which were conducted on operational networks during early morning periods of light network activity, the low amounts of live traffic on the networks and the variable processing delays of the hosts sufficed to add variability to the experiments.

Behavior of Several TCP Variants

We plot the results of these experiments in Figure 4.13. In all of our figures, throughput is defined as "application-level" throughput. For values of RTT less than 100 ms, the performance is relatively high for all four variants. However, for GEO delays (600 ms) and for LEO delays greater than 100 ms, the difference in performance for different TCP implementations is quite evident. By analyzing packet traces in both the simulations and the experiments, we determined that the main distinction between the implementations was in their behavior immediately upon leaving the slow start phase of congestion avoidance. It is critical that TCP transition from slow start to congestion avoidance in a smooth manner, with a congestion window close to the bandwidth-delay product of the path. We found the performance of SACK-NewReno congestion avoidance to be the best; in this case, when a slow start overshoot occurs, the protocol cuts its window in half once and smoothly moves to congestion avoidance after recovering all losses. There is little penalty for using a high-bandwidth, high-latency GEO satellite link in this case. When SACK was used without NewReno enhancements (SACK-Reno), we observed that the slow start termination, which is characterized

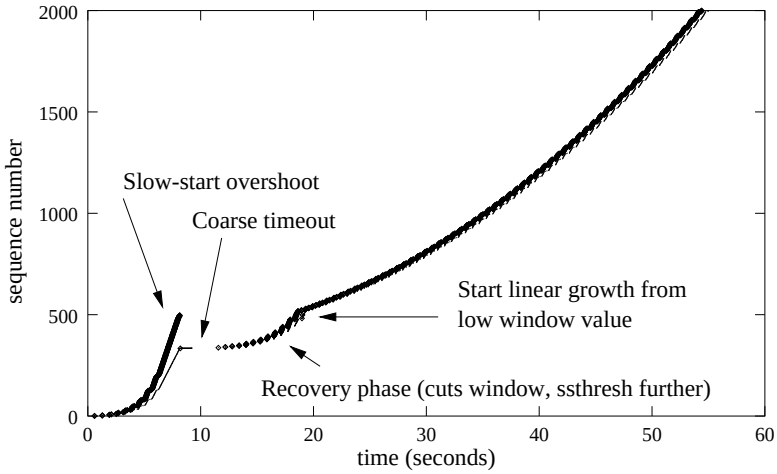

Figure 4.14: Typical performance, using a standard BSD TCP (Reno) implementation, of a large file transfer over a GEO satellite channel.

by several bursts of packet losses, resulted in the implementation cutting its congestion window in half several times, rather than just once. As a result, TCP was forced to rebuild its window linearly from a very low value. The performance of NewReno without SACK was similar but for a different reason. In this case, the slow start overshoot resulted in similar bursty patterns of losses, but since NewReno, unlike SACK, can only recover one loss per RTT, it spent a large portion of time recovering from the slow start losses. Finally, TCP Reno rarely avoided a retransmission timeout and multiple reductions in its window after the first slow start, resulting again in slow window growth.

A closer look at the behavior of these different TCP variants is informative. Figure 4.14 illustrates a "time-sequence" plot of an individual connection– the initial 60 seconds of a large file transfer using an unmodified BSD/OS UNIX TCP implementation (TCP Reno without SACK) over the topology illustrated in Figure 4.11. Two plots are overlaid– the evolution of the sender's sequence number, and the evolution of the acknowledgments received (the trace of the sender's sequence number generally lies to the left of the acknowledgment trace and is marked by larger points). The connection initially starts in slow-start, and although the connection takes several seconds to make noticeable progress, within the first 10 seconds the connection has already overshot by a wide margin the capacity of a router along the path, resulting in many packet drops (not necessarily contiguous in the sequence space). The implementation performs fast retransmission, but since many packet losses have occurred, the implementation invariably is forced to recover with a coarse timeout because it does not interpret the arrival of a partial acknowledgment (that is, an acknowledgment that does not cover all of the data that was outstanding at the time of the retransmission) as a sign that the next unacknowledged packet is missing. The timeout cuts the window to one segment and the slow start threshold in half; normally this would allow the sender to rapidly ramp up after the timeout to half of its previous window (that caused congestion). However, because the receiver's buffer has many out-of-order packets (and holes to fill), as the post-timeout

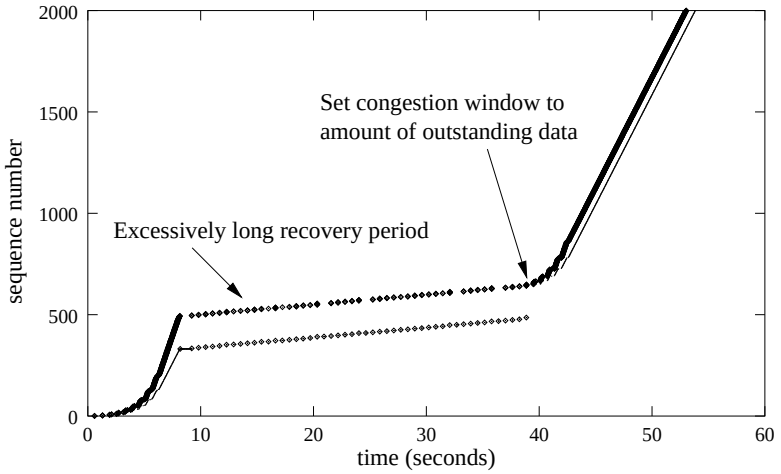

Figure 4.15: Correct NewReno behavior, using a modified BSD TCP implementation, of a large file transfer over a GEO satellite channel.

TCP sender starts to send more data, it can receive a number of duplicate acknowledgments that push it in and out of fast retransmission again (these are sometimes called "false fast retransmissions" [60]), each time cutting the window and slow start threshold in half. The result is, by the time that the sender has recovered from all of the original losses, it has a very low congestion window and slow start threshold value, and is forced to build its window linearly from a very low value, resulting in poor throughput.

TCP NewReno was devised to correct this oversight in the TCP Reno implementation; it defines a "recovery phase" that ends when all of the packets that were outstanding when the first loss was detected are acknowledged. Figure 4.15 illustrates the typical performance of this algorithm. No timeouts occur during the recovery and the window is not reduced multiple times for the same burst loss. However, since the recovery takes one round trip time for each gap in the sequence space to be recovered, the result is a TCP connection that takes over half a minute to recover from a single burst of losses. For this reason, as pointed out by Floyd [42], it may be beneficial to prevent this behavior from occurring by forcing TCP to take a timeout if it requires too many round trips to recover. Second, in our experiments with this algorithm as specified by [60] and [36], we noticed undesirable behavior that occurred at the end of the recovery phase. This behavior (immediate reentry into a burst loss situation) can be seen in Figure 4.16, and it is due to a burst of packets that can occur at the end of recovery. This burst occurs if, during window inflation of the recovery phase, the transmission of new segments was constrained by the receiver's offered window (because the "window inflation" step of TCP Reno can result in very large artificial windows being generated). As a result, when the hole is plugged in the reassembly buffer and the TCP sender resets its congestion window upon receipt of the acknowledgment, it is eligible to immediately send many segments. The solution to this problem, as shown in Figure 4.15, is to constrain the congestion window at the end of recovery to be no larger than the amount of outstanding data at that time (plus one segment, to allow a new transmission). This proposal was first described in [76].

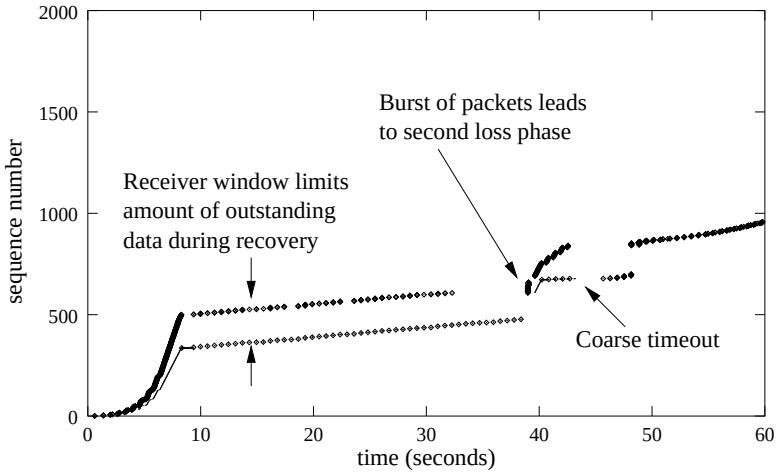

Figure 4.16: Incorrect NewReno behavior, using a modified BSD TCP implementation, of a large file transfer over a GEO satellite channel. This behavior is due to a burst of packets transmitted at the end of recovery.

The best behavior is obtained by combining both the SACK and NewReno algorithms, as illustrated by Figure 4.17. In this case, TCP recovers very rapidly from bursty losses because the extra information present in the SACK option gives the TCP sender a more complete picture of what is missing. Figure 4.18 illustrates this recovery in more detail; the burst of losses is recovered in less than two seconds (that it requires more than a round trip delay is due to queueing delays that have built up), and a large window is preserved for the subsequent connection to use in linear growth phase.

Finally, we illustrate in Figure 4.19 the close correspondence between simulation and experimental results for file transfers. In the remainder of this section, we present only our experimental results since our simulation results were generally in close agreement. The TCP implementations of the *ns* simulator are very realistic, to the point that bugs found in common implementations can also be enabled in our simulations.

Effect of a Competing Connection

The above experiments are appropriate to model connections entirely within a satellite subnetwork, but do not accurately portray conditions found when using the satellite network to access sites on the wired Internet, where competition for bandwidth from many different connections (with shorter round trip delays) can lead to network congestion and unfairness in bandwidth allocation. For our next experiments, we added a single, large-window persistent connection from a background source to a background sink in the same direction as the foreground file transfer. In our topology, this caused the first router in the network to occasionally become congested. Note that this background connection does not traverse any portion of our emulated satellite subnet. The results in this case are strikingly different. It only takes one low delay (in this case, 20 ms RTT) connection to drastically reduce the achievable throughput for SACK-NewReno, as shown in Figure 4.20. This

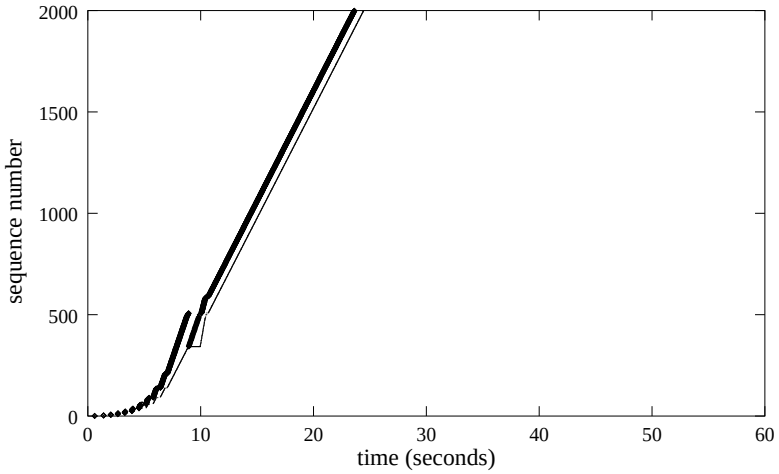

Figure 4.17: Correct SACK behavior, using a modified BSD TCP implementation (including NewReno loss recovery), of a large file transfer over a GEO satellite channel.

is the TCP fairness problem identified earlier in this chapter. TCP's fairness properties can be the first-order determinant of how well a large-window satellite TCP connection can do in the wide-area Internet. Even though the satellite connection was successful in avoiding timeouts in almost all of the transfers, the window reductions due to recurring fast retransmits substantially reduced the throughput. The throughput is also much more variable under these conditions, as represented by the error bars. The main problem is that the connection with the long RTT is too sluggish to rebuild its window and push data through the congested queue before it takes another loss.

In summary, we observed that TCP SACK with NewReno congestion avoidance is able to sustain throughputs at close to the bottleneck link rate even for GEO-like delays. This is because TCP is able to amortize the low throughtput of the initial window build across a longer period of high throughput. However, our data illustrates that the use of SACK alone is not sufficient to enable high performance. Specifically, NewReno helps to avoid coarse timeouts and multiple window reductions, while SACK accelerates the loss recovery phase. Specific details of our SACK-NewReno implementation can be found in Appendix A. Finally, the result we would like to emphasize, which agrees with our analysis in Section 4.1, is that it only takes very moderate levels of congestion in the wide-area Internet to drastically impair the performance of even well-configured TCP connections.

4.2.4 Performance for Web Transfers

Besides file transfers, most of the rest of the TCP traffic in the Internet is driven by Web transfers. Such connections are very different from file transfers. Typically, an Web client issues a small request to a server for an HTML (HyperText Markup Language) page. The server sends the initial page to the client on this first connection. Thereafter, the client launches a number of TCP connections to fetch images that fill out the requested page or to obtain different pages. Each item

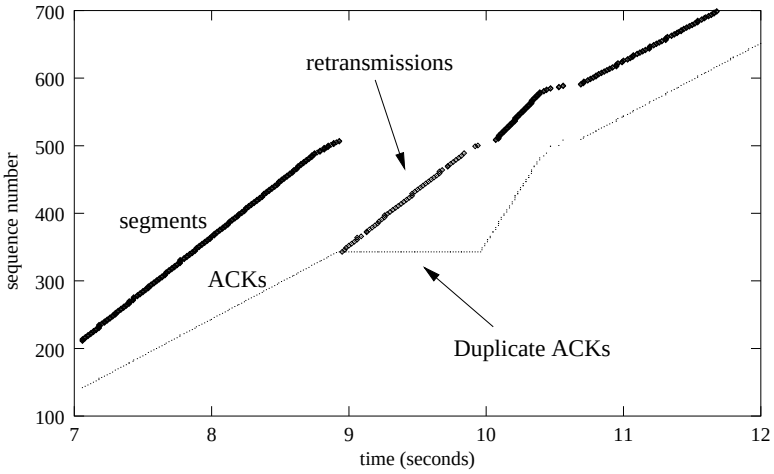

Figure 4.18: Closeup of rapid SACK recovery of multiple losses.

on the page requires a separate connection.[6] Many common Web browsers allow a user to operate multiple (typically, four) TCP connections in parallel to fetch different image objects. Basically, the data transfer model is "client request, server response."

Using standard TCP, any connection requires a minimum of two RTTs until the client receives the requested data (the first RTT establishes the connection, and the second one is for data transfer). As the RTT increases, the RTT can become the dominant portion of the overall user-perceived latency, particularly since average Web server response times are much smaller than one second [51]. Two mechanisms described in Section 3 attempt to alleviate the latency effects of TCP for short connections. The first, T/TCP, does away with the initial handshake (RTT) of the connection. The second, 4KSS, allows the TCP server to send up to 4380 bytes in the initial burst of data. If the size of the transfer is no more than 4380 bytes, the transfer can complete in one RTT. By using some simple analysis, we can quantify the beneficial effects that these TCP mechanisms have on the user-perceived latency.

Figure 4.21, adapted from a similar figure in [57], illustrates the latency in a hypothetical three segment reply using standard TCP. We make the following assumptions:

- We do not model server response times or segment transmission times. We assume an environment in which the RTT is the dominant latency in the transfer.[7] Server response times and segment transmission delays are a constant offset to the latencies we calculate; i.e., the same offset must be added no matter what version of TCP we are considering.

- We assume no packet losses and a fixed RTT. Therefore, these latencies are the best case.

- We do not model some of the bugs that have appeared in early HTTP implementations and that

[6] We will discuss shortly a modification to this approach, known as Persistent-HTTP (P-HTTP), which reuses the same TCP connection for multiple items.

[7] This is not always true in practice. Even for fast links, server responses can take several seconds, but on average, the server response time is much less than a second [51].

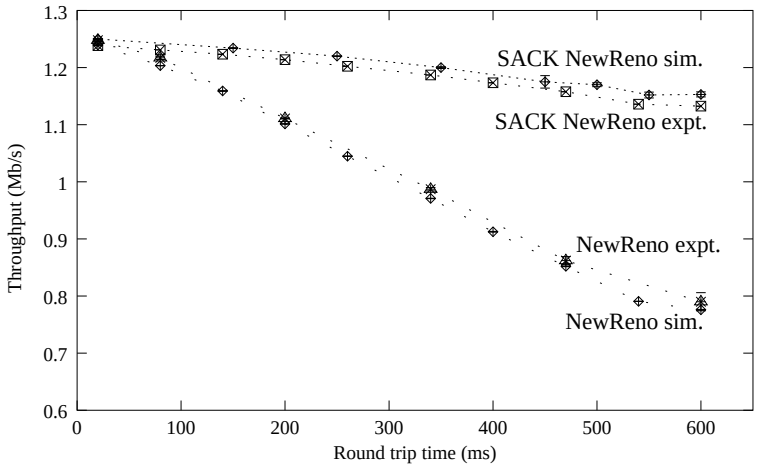

Figure 4.19: Agreement between simulation and experimental results for TCP SACK and TCP NewReno.

are discussed in [57], under the assumption that they will gradually disappear. For example, one quite prevalent bug allows the connection to start with an initial congestion window of two segments [128].

With these assumptions in mind, consider Figure 4.21, in which dashed lines denote control packets and solid lines indicate data packets. The first RTT is consumed by a SYN exchange, after which the client issues an HTTP GET request. Upon receiving and responding to this request, the server at this point has a congestion window of one segment. Assuming that the TCP implementation implements delayed acknowledgments (delayed ACKs) of up to 200 ms [127], the client on average will acknowledge this data after 100 ms. Upon receiving the acknowledgment, the congestion window grows to 2, and the server sends the second and third segments, followed by a FIN, which closes its half of the connection. The client must close its own half of the connection, but we do not model this delay since it does not contribute to user-perceived latency. Therefore, the total amount of TCP-related latency is 3 RTTs + 100 ms in this case. Using either T/TCP or 4KSS would reduce the latency to 2 RTTs, and using both mechanisms would reduce it to a single RTT.

We used HTTP traces to compute probability mass functions (pmfs) for the number of bytes transferred per HTTP connection. We then computed the average TCP latency for all of these file sizes, based on a simple analysis of how the congestion window builds over time. Because some transfers were very long, we eliminated those over 100 segments (only 2-4% of the data set, in general). For these cases, it is more realistic to consider them as large file transfers. Our trace data was gathered from two different user populations. The first, collected by Mah in 1995 [79], comes from a well connected Berkeley subnet. The second set, collected by Gribble in 1997 [51], comes from Berkeley residential usage over dial-up modems. By using this trace data with our model, we estimated the minimum, median, and mean latency effects of TCP on user-perceived latency. For GEO networks, we modeled the RTT as a fixed 600 ms, and for LEO networks we assumed a RTT of 80 ms. To verify the analytical results, we also performed measurements using similar pmfs to

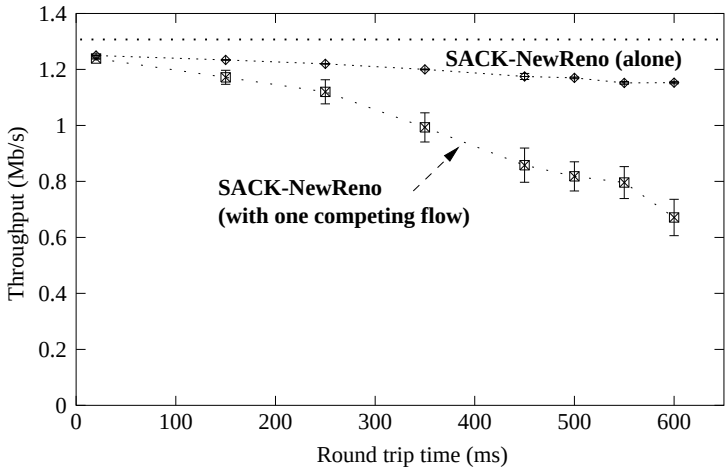

Figure 4.20: The effect of a single competing short-delay connection on the satellite connection's throughput. The competing connection was a persistent file transfer using TCP SACK NewReno with a nominal 20 ms RTT between a background source and sink in the experimental topology.

drive a TCP traffic generator in our experimental topology, and we recorded the latency experienced along with the file size for each file transfer. For the experiments, we did not cull the large transfers from our trace data. The experimental setup captured the effects of not only the propagation delay but also the processing delays in real end systems.

In Table 4.2, we present the results from an analysis of the data set provided by Mah [79]. The first three columns of data list the minimum, median, and mean TCP transfer times required, according to the analysis of the trace file and assuming a maximum segment size of 1500 bytes. These values were calculated by first determining the TCP related latency for a connection of a given size, and then by weighting these latencies according to the pmfs derived from the trace data. The fourth column lists experimental results corresponding to this data set. These values are the mean (and 95% confidence interval) of 1000 independent transfers, in which the size of the transfer was generated randomly according to the pmfs drawn from the trace data. The last four columns are similar to the first four, except for the use of a maximum segment size of 500 bytes. This data indicates that the use of either T/TCP or TCP with 4KSS improves mean latency by a small amount, but the combination of both options yields an improvement by a factor of two to three. The relative improvement is similar whether GEO or LEO networks are assumed (because the analysis is based on RTT). Because the mean latencies using the assumed LEO network are already rather small, the improvements due to TCP optimizations are less likely to be perceived by users. The data set provided by Gribble [51] contained slightly larger transfers, on average, but the same trends in TCP latency were present.

Finally, the most recent version of the HTTP specification (version 1.1 [37]) recommends that servers and clients adopt the persistent connection and pipelining techniques known as "persistent-HTTP" (P-HTTP) [102]. Rather than using separate TCP connections for each image on a page, P-HTTP allows for a single TCP connection between client and server to be reused for

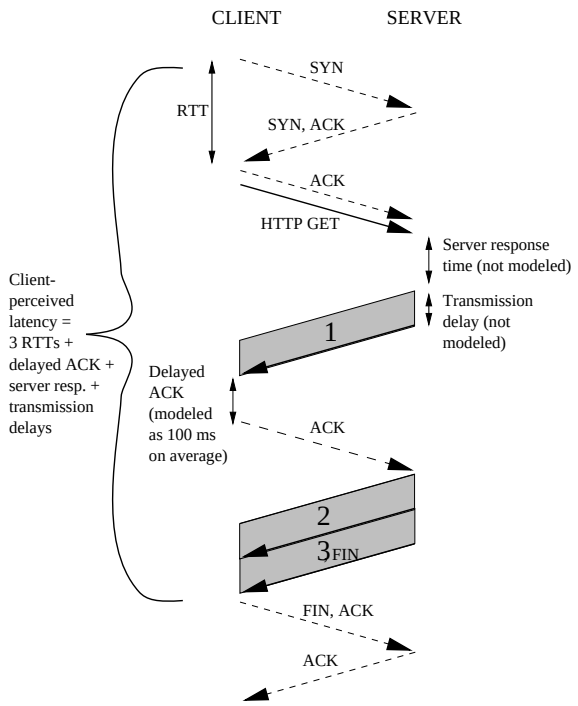

Figure 4.21: TCP latency of a 3 segment server reply using standard TCP.

multiple objects. The shift to P-HTTP offers a tradeoff in performance for satellite connections. On the one hand, P-HTTP is potentially much more bit-efficient than HTTP with standard TCP, because connections are not set up and torn down as frequently (the connection establishment costs are identical to those of T/TCP [57]). However, in terms of latency, the use of T/TCP and multiple, concurrent connections may yield faster Web page loads under some scenarios. The capability of many Web browsers to support multiple, concurrent connections is an example of a general technique known as "striping," which has been a strategy for transport protocol improvement known to satellite network operators for some time, and which has most recently been studied in the context of FTP [6]. Because TCP and HTTP optimizations such as T/TCP, and TCP with 4KSS do not yield major performance improvements for most users of the Internet [57], it is unclear whether they will see deployment. In fact, Padmanabhan recently studied the potential benefit of not using P-HTTP but instead reverting back to multiple, concurrent TCP connections that share congestion window and other state information [99].

In summary, for connections using GEO satellite links, TCP optimizations such as T/TCP and 4KSS, especially when used together, can yield a reduction of two to three times in in user-perceived latency and can also reduce the bandwidth overhead of HTTP connections. For LEO satellite links, optimizations to reduce the number of unnecessary control packets are desirable, but optimizations to reduce latency will not have as perceptible of an effect for users because propagation delays are smaller. However, since such optimizations benefit only a small user community, it is possible that they will not see widespread deployment.

	Geostationary orbit (600 ms RTT)							
	1500 byte segments				500 byte segments			
	minimum	median	mean	expt. mean	minimum	median	mean	expt. mean
Standard TCP	1.2	1.9	1.9	2.0 (0.1)	1.2	2.5	2.5	2.6 (0.1)
T/TCP	0.6	1.2	1.2	1.4 (0.1)	0.6	1.8	1.7	2.0 (0.1)
TCP "4KSS"	1.2	1.2	1.4	1.6 (0.1)	1.2	1.8	1.7	1.9 (0.1)
T/TCP "4KSS"	0.6	0.6	0.8	1.0 (0.1)	0.6	1.2	1.1	1.3 (0.1)
	Low-earth orbit (80 ms RTT)							
Standard TCP	0.16	0.34	0.31	0.37 (0.02)	0.16	0.42	0.42	0.55 (0.02)
T/TCP	0.08	0.16	0.17	0.28 (0.02)	0.08	0.24	0.25	0.47 (0.02)
TCP "4KSS"	0.16	0.16	0.18	0.25 (0.01)	0.16	0.24	0.23	0.31 (0.01)
T/TCP "4KSS"	0.08	0.08	0.10	0.16 (0.01)	0.08	0.16	0.15	0.23 (0.01)

Table 4.2: TCP latency effects on HTTP transfers for GEO and LEO satellite connections. Trace data is taken from [27]. All latencies are in seconds. For the experimental results, 95% confidence intervals are shown in parentheses.

4.3 Split TCP Connections

Although TCP can work well over even GEO satellite links under certain conditions, we have illustrated that there are cases for which even the best end-to-end modifications cannot ensure good performance. Furthermore, in an actual network with a heterogeneous user population, users and servers cannot all be expected to be running satellite-optimized versions of TCP. This has led to the practice of "splitting" transport connections. This concept is not new; satellite operators have deployed protocol converters for many years. In this section, we describe how TCP connections may be split at a satellite gateway, identify some drawbacks to split connections, and quantify how much improvement can be obtained.

4.3.1 Split Connection Approaches

The idea behind split connections is to shield high-latency or lossy network segments from the rest of the network, in a manner transparent to applications. TCP connections may be split in a number of ways. Figure 4.22 illustrates the most general case, in which a *gateway* is inserted on the link between the satellite terminal equipment and the terrestrial network. On the user side, the gateway may be integrated with the user terminal, or there may be no gateway at all. The goal is for end users to be unaware of the presence of an intermediate agent, other than improved performance. From the perspective of the host in the wide-area Internet, it is communicating with a well-connected host with a much shorter latency. Over the satellite link, a satellite-optimized transport protocol can be used.

TCP may be split in the following ways:

- **TCP spoofing** In this approach, the gateway on the network side of the connection prematurely acknowledges data destined for the satellite host, to speed up the sender's data transmission [146]. It then suppresses the true acknowledgment stream from the satellite host, and

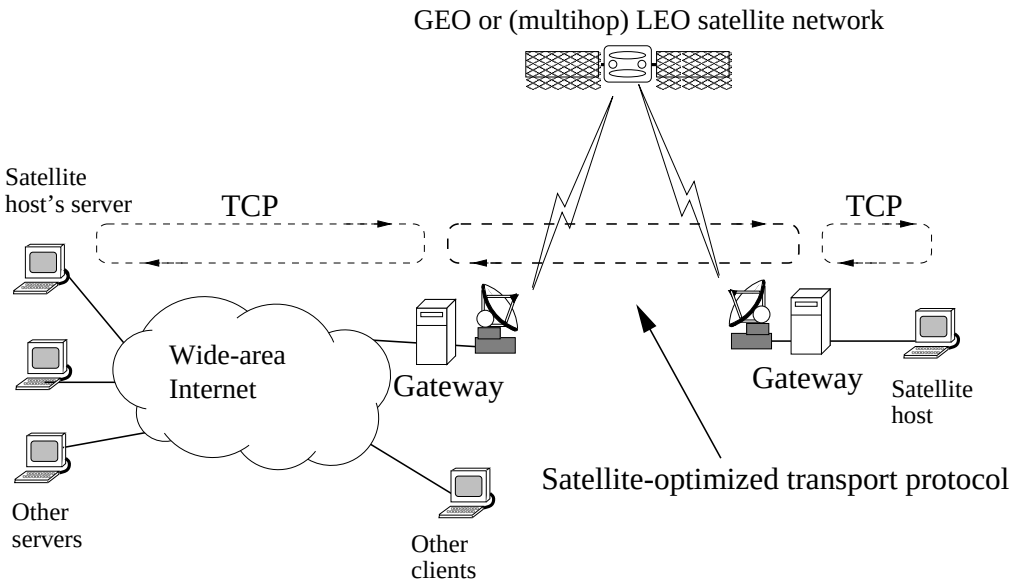

Figure 4.22: Future satellite networking topology in which a satellite-based host communicates with a server in the Internet through a satellite protocol gateway.

takes responsibility for resending any missing data. As long as the traffic is primarily unidirectional, TCP datagrams are passed through the gateway without alteration. In the reverse direction, the same strategy is followed. No changes are needed at the satellite client.

- **TCP splitting** Instead of spoofing, the connection may be fully split at the gateway on the network side, and a second TCP connection may be used from the satellite gateway to the satellite host. Logically, there is not much difference between this approach and spoofing, except that the gateway may try to run TCP options that are not supported by the terrestrial server. Modern firewall implementations often perform a type of TCP splitting (such as sequence number remapping) for security reasons.

- **Web caching** If satellite-based Web users connect to a Web cache within the satellite network, the cache is effectively splitting any TCP connection for requests that result in a cache miss. Therefore, Web caching not only can reduce the latency for users in fetching data from the Web, it has the side benefit of splitting the transport connection for cache misses.

Furthermore, when the TCP connection is fully split at a gateway or cache, it is possible to use an alternative protocol for the satellite portion of the connection. While this requires the use of a satellite gateway or modified end-system software on the satellite host's side, this approach may provide better performance by improving on TCP's performance in ways not easily achieved by remaining backward compatible with existing implementations. Set-top boxes or other user terminal equipment may provide a natural point for the implementation of protocol conversion (back to TCP, if necessary) on the satellite host's side of the connection.

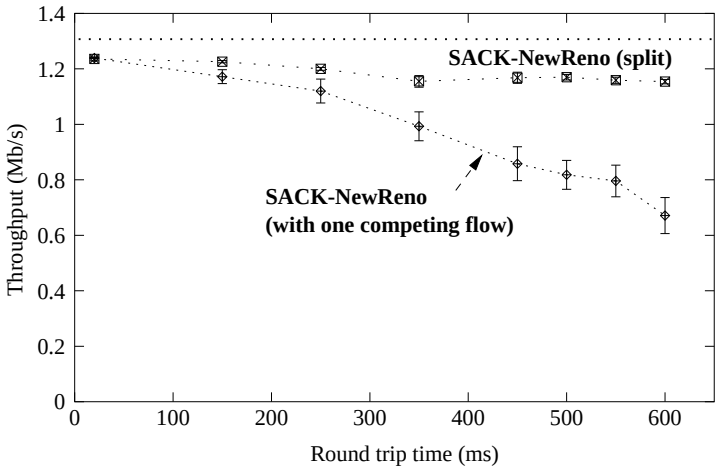

Figure 4.23: Forward throughput performance of split TCP in the presence of a short-delay competing connection. TCP SACK NewReno with large windows was used on both connection portions.

In all three approaches, the amount of per-connection buffering required at the gateway is roughly 2-3 times the bandwidth-delay product of the satellite link or the Internet path, whichever is smaller. The computing resources required to support a large set of users (approximately 200-500 KB of memory per active connection, plus processing) can form a significant portion of the hardware requirements of a satellite Internet gateway. In addition, although persistent-HTTP connections will reduce the number of connections that need to be set up and torn down, they will also drastically lower the duty cycle of each TCP connection, requiring the gateway to keep resources allocated for idle connections. However, it is important to emphasize that if Web caches or other proxies are already part of the satellite network architecture, there would be no need for extra equipment to support transport-level gateways.

Besides the resource consumption noted above, split connections are not without other hazards. First, from an architectural standpoint, a split TCP connection that is not explicitly associated with a proxy or a cache breaks the end-to-end semantics of the transport layer. Although approaches for TCP improvement over local area wireless links, such as Berkeley's "snoop" protocol [10] and "mobile TCP" [21], can preserve end-to-end semantics, it is more difficult to do so in the satellite environment because combating the fairness problem relies on early acknowledgment of data. However, steps can be taken to ensure that the connection does not close normally unless all data has been received; for example, the gateways can allow the FIN segment of TCP to pass end-to-end. Furthermore, higher layer protocols typically have mechanisms to restart a transport connection if it prematurely fails. Second, gateways introduce a single point of failure within the network, and require all traffic for a given connection to be routed through them (i.e., there can be no alternate packet routing). Third, protocol conversion gateways are ineffective if IP-level encryption and authentication protocols are operating on a link, although they can still function normally if the encryption and authentication is performed at the transport layer. In the case of IP-level security, the transport gateway must be included as part of the "trust infrastructure" to operate. Typically,

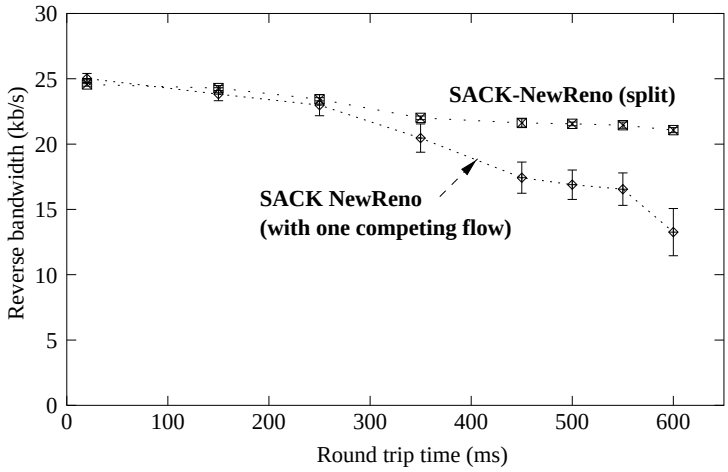

Figure 4.24: Reverse channel utilization of split TCP in the presence of a short-delay competing connection. TCP SACK NewReno with large windows was used on both connection portions.

however, if a satellite network is used to provide "last-mile" access to a large diverse set of users, transport-level security protocols will be used instead of IP-level security; in this case, protocol gateways can operate correctly.

4.3.2 Split Connection Performance

In Figure 4.23, we illustrate the performance gains achievable when the TCP connection is split at the gateway between the satellite network and the Internet, under the same conditions as shown in Figure 4.20 (a competing short delay connection in the Internet). We replotted the relevant data from Figure 4.20 for comparison. Note that the presence of the gateway allows the split connection to compete for bandwidth in the wide area and obtain its fair share. However, as shown in Figure 4.24, the reverse channel usage required for this TCP connection is roughly 20 Kb/s. This usage scales linearly with the forward throughput, and for 1000 byte segments, is roughly 2% of the forward throughput achieved. For bandwidth-constrained reverse channels as will be the case in most satellite systems, this sets an upper bound on the forward throughput achievable if TCP relies on a stream of acknowledgments to clock out new data. This suggests that it would be useful either to make modifications to TCP to reduce its reverse channel usage (such as using modifications to handle TCP asymmetry [11]) or to use a protocol over the satellite portion of the connection that uses less bandwidth. We investigate the latter possibility in the next chapter.

4.4 Summary

In this chapter, we have investigated the performance of IP-compatible transport protocols over satellite links from several perspectives. Our main results are as follows:

i) We observed little degradation in TCP performance for connections with RTTs in the

range of future LEO systems (40-200 ms), although we did not investigate potential problems due to large RTT variations. However, maintaining good TCP performance over GEO latencies (or long LEO paths) is challenging.

ii) While we found that the Constant-Rate (CR) policy could improve TCP fairness dramatically, we faced two practical difficulties that would likely prevent universal deployment of this scheme in its current form. First, the proper selection of a constant is dependent upon the network topology and the number of peer connections and is therefore difficult to determine in a distributed manner. Second, the fairness benefits of the CR policy can be confounded by competing connections using standard congestion avoidance, thereby making it disadvantageous to deploy CR in an existing heterogeneous environment.

iii) When we instead made only certain long RTT connections slightly more aggressive, we were always able to improve network fairness while keeping bottleneck link utilization relatively constant by using an increase-by-K (IBK) policy. Interestingly, the effects on other unmodified connections that were sharing the bottleneck link were similar to what they would have experienced had the modified connection actually been a connection with a shorter RTT.

iv) If the right TCP options are used and congestion is light, TCP can work well for large file transfers even over GEO links. In particular, in our large file transfer experiments with TCP SACK plus NewReno congestion control, average throughput decreased by no more than 10% when the RTT was increased from 20 ms to 600 ms. However, we showed that even low levels of competition from short delay flows (in the form of cross-traffic in the wide-area Internet) significantly degrades the satellite connection's performance.

v) Concerning the latency due to HTTP exchanges, we found that the use of both T/TCP and modified slow start performed much better than either option used separately, and could cut the average TCP-related latency by a factor of two to three for GEO links.

vi) We showed that the performance problems due to mis-configured TCP or network congestion can be alleviated by splitting the TCP connection at a gateway within the satellite subnetwork. Even with congestion in the wide-area Internet, the end-to-end connection is still able to maintain high throughput.

TCP has proven to be a very robust protocol in a variety of network environments. However, this chapter has illustrated that obtaining good performance using standard end-to-end TCP connections is very challenging in a GEO satellite environment. For file transfers, the best performance results that we obtained were based on splitting the connection at a gateway, where the long round trip delay of the satellite portion of the path can be isolated from the portion of the connection that traverses the Internet. For short transactions, we found that the best performance requires TCP enhancements (T/TCP, 4KSS) that are not implemented in the current Internet– again leading us to consider split connection solutions. Given that one decides to split TCP connections at a satellite gateway, it is natural to ask what protocol (either modified TCP or an entirely new protocol) should be used over the satellite portion of the split connection. In the next chapter, we attempt to improve on our performance results even further by designing a satellite-optimized transport protocol that outperforms even the split TCP configuration described above in Section 4.3.